|CURIEUX|

ACADEMIC REVIEW

Physical Sciences

2026 Issue Part 1

Editing Staff

Chief-In-Editor

George Hasnah

Lead Editor

Amanuel Geremew

Ashley Du

Assisting Editors

Jing Yan Huang

Nidhi Potineni

Sophia Kong

Tony Yuan

Ridit Aggarwal

Manu Sharma

Rishi Gadikota

Eric Wang

Table Of Contents

Page 6: Nanotechnology in Food Packaging: A Breakthrough for Shelf Life or a Catalyst for Biological Risk? By Jaiaghasthya Ashokkumar

Page 18: A Bioinformatics-Guided Study of Moringa–Plant Polymer Coagulants for Sustainable Microplastic Removal from Water By Neil Lewis[1] and Nirupma Singh[2]

Page 29: Comparative Evaluation of the Physiochemical Properties of Gluten and Gluten-Free Blondies By Isabella Jacobson

Page 49: Investigating The Effects of Centroid Motion on Exoplanet Detection Using TESS Light Curve Data By Advik Sharma

Page 64: A Brief Exploration of Black Holes in the Early Universe By Reuben Saunders

Page 69: How Does the Kauffman Bracket with its Writhe Correction Prove that the Jones Polynomial is a Well-defined Knot Invariant? By Zhiheng Luo

Page 85: Orogeny of Mount Gyoja: Shaped through Opening Processes By Néo Petit

Page 106: Anti-Knocking Agents in Gasoline: A Comprehensive Review of Mechanisms, Performance, and Environmental Impact By Dhairya Vij

Page 119: The Effects of pH Level on the Decomposition of Xenobiotic Vinyl Compounds and the Resulting Health and Environmental Issues By Stella Dybczakova

Page 126: Running on Empty? : Galactic Quenching Mechanisms and Gas Retention through Morphology, Mass, and Environment By Aarav Verma

Page 142: Beginner's Guide to the Hubble Tension: The Basics By Kento Noda

Page 150: Does ZEV Adoption Improve Regional Air Quality? Empirical Evidence from Los Angeles County (2010–2025) By Yinuo Jiang and Kaixi Chen

Page 165: Evaluating the Economic Scalability and Energy Efficiency of Magnetic Confinement Strategies Adopted by Emerging Fusion Startups for Future Commercial Fusion Power Plants By Anahita Dubey

Page 191: Evaluating How MATLAB in Python Can Be Used to Plot Planck's 'Black Body Radiation' Spectrum By Akriti Srivastava

Page 196: A Comparison Between Direct Heating and Temperature Control Using a Water Bath By Dahyuck Im

Page 205: Beyond Chance and Necessity: Cosmic Fine-Tuning at the Boundary of Cosmology and Teleology By Camila Richardson

Page 216: A Review of Sustainable Battery Design: Alternatives to Lithium-Ion Batteries: Case Studies of Redox, Na^+, Mg^{2+} Ion Based Battery By Shloke Kusoorkar, Lawrence T. Tom

Page 238: Agrivoltaics in California: Balancing Water, Energy, and Farmland in the Central Valley By Angela Jin

Page 258: Comparison of Novel Aluminum-Free Deodorants and a Traditional Mineral-Based Deodorant By Joe Fang

Page 267: Investigating the Efficiency of Melanin as an Active Layer Absorber in Perovskite-Inspired Solar Cell Architecture via SCAPS-1D Simulation By Hamza Ahmed

Page 280: Predicting Weather Patterns for Rochester, New York By Wohaan Dissanayake

Page 286: Topology in Electromagnetic Confinement: Literature Review and Conceptual Synthesis By Hassan Shahid

Page 305: Spatio-Temporal Analysis of Urban Air Pollution Dynamics in Delhi-NCR: Hotspot Identification, Seasonal Trends, and Inter-Pollutant Correlations By Aadish Garg

Page 321: A Review of Environmental Changes Affecting Coastal Winds Near Nantucket and Effects on Sailing By Sammy Laliberte

Page 328: How Does the Rotation Curve of a Spiral Galaxy Deviate from Keplerian Predictions, and What Does This Imply About the Distribution of Dark Matter? By Agastya Nambiar

Page 346: Steel Decarbonisation and Cement Sustainability: A Literature Review of Slag-Based Industrial Symbiosis By Leon Langford

Page 371: Influential Factors of Effective Encapsulation Design in Fragrance Delivery By Liana Kim

Page 379: Drag Force on a Wigley Hull: Equation vs. Simulation Across Velocities and Drafts By Rithikk Vimal

Page 389: Timescape Cosmology: An Alternative to the ΛCDM Standard Cosmological Model? By Bogdan Stoiljkovic

Page 405: A Review Analysis of the Feasibility of Implementing RGBW Sensors in Cameras By Kiaan Shah

Page 416: A Mathematical Analysis of How Harmonic Frequencies Explain Differences in Musical Timbre Across Instruments By Guanhua Yang

Page 429: Using Ising Model to Explain Demagnetization Phase Transition of Magnet By Zihan Zhou

Page 438: How do Wavelength-Dependent Transit Depths Reveal the Presence of Various Chemical Species in the WASP-121b Atmosphere? By Aadya Lal

Page 447: Gravitational-Wave Memory Effect and How Its Detection Could Reveal About Spacetime and Gravity By Pascale Gani Sutantio

Page 455: A Trigonometric Model of a Vertically Oscillating Arrow Passing Through a Small Hole By Junhyeok Sung

Page 464: PM2.5 Trends in Almaty (2017–May 2026) in the Context of the CHP-2 Coal-to-Gas Conversion Project By Arailym Oral

Page 472: Why Are Three-Point Shots Mathematically Advantageous in the NBA? By Taehoon Kim

Page 484: Sustaining the Ski Slopes: Snow Preservation through Biodegradable Geotextiles By Connor Mason

Page 500: Testing the 70-30 Rule: A Simulation-Based Evaluation of Aggressive and Conservative Play in Junior Girls' Golf By Kexin Zhang

Page 516: Heteroscedastic Gaussian Process Regression for Out-of-Domain Neutrino-Nucleus Cross-Section Prediction By Atharva Vignesh, Bhavya Darji, Yug Dhami, and Rohan Kumar

Nanotechnology in Food Packaging: A Breakthrough for Shelf Life or a Catalyst for Biological Risk? By Jaiaghasthya Ashokkumar

Abstract

Nanotechnology is a recent development that has shown great potential for food packaging by enabling better preservation, longer shelf life, and reduced food waste. Incorporating nanomaterials like silver, zinc oxide, titanium dioxide, nanocellulosic compounds, and nanosensors in packaging materials, researchers have created new materials with high antimicrobial activity, improved barrier properties, and intelligent spoilage-detecting capabilities. This paper reviews the principal application areas of nanotechnology in food packaging and highlights key issues, including nanoparticle migration, human health, and environmental consequences, as well as regulatory frameworks in the US and EU. According to the literature, food packaging with nanoparticles helps enhance the quality of the product and reduce the spoilage, but there are still a number of issues that remain, such as the effect of nanoparticles on humans in the long-run, movement of nanoparticles from packaging into the food, and the effects after disposal. In addition, different regulatory systems between countries demonstrate further the necessity to have regular safety assessments and unified procedures in testing nanomaterials designed for food-contact use. In a more general sense, nanotechnology could be a novel approach to making food safety safer and more economical for resources, but it would take additional toxicological studies, migration investigations, and uniform regulations across borders for all the countries to carry out its risk-free and safe deployment.

Keywords

Nanotechnology; Food Packaging; Food Safety; Nanoparticles; Antimicrobial Packaging

Introduction

Nanotechnology is the science and engineering of ultrafine materials at approximately 1 to 100 nanometers, and this technology is one of the most impactful scientific advancements of the twenty-first century. A change occurs at this tiny scale, such that materials often possess some qualities which they do not exhibit in a bulk state, e. g., they can be stronger, more chemically reactive, or able to prevent microbial growth more effectively. The different behaviors of materials at different lengths have resulted in different applications in fields like medicine, electronics, and food science. One area in which nano food research is expected to show significant progress is packaging. Since packaging protects the food from external elements Make sure to ensure its quality and integrity for a longer time. That's why even the smallest enhancement in its performance would be able to not only delay spoiling but also decrease the amount of food wasted. Unfortunately, the same features (e. g., nanoparticles' minute size and intense activity) that render them so attractive for food packaging might provoke concerns, such as whether the particles will still stay there or what the implications would be for their use at such levels for humans and ecosystems.

Nanotechnology is, because of this, a powerful illustration of how a relatively controversial scientific development may at the same time offer a real solution and introduce a new set of risks. It is a good example of where the same material that can solve a difficult food spoilage problem is also asking the public and the governmental bodies to give their trust in a very young category of substances, the impact and behavior of which are only gradually being studied and understood. Because of this, this overview considers nanotechnology as a means of food packaging while discussing its applications, potential health and environmental concerns to the human, and the regulatory setups controlling it.

Literature Search Method

A literature review was conducted by searching peer-reviewed articles and reports from scientific databases such as Google Scholar, PubMed, ScienceDirect, SpringerLink, MDPI, and official publications from regulatory agencies, e.g., U.S. Food and Drug Administration (FDA), United States Department of Agriculture (USDA), Environmental Protection Agency (EPA), and European Food Safety Authority (EFSA). Research was carried out by using mix and match sets of words: nanotechnology, food packaging, silver nanoparticles, zinc oxide, titanium dioxide, nanocellulose nanosensors, food safety, food waste, nanoparticle migration. Main effort has been concentrated on peer-reviewed publications and government reports dated between 2015 and 2027. The most groundbreaking research published well before this timeframe was cited only when a background or a historical context was needed. The choice of sources was largely dictated by themes that cover nanotechnology uses in food packaging, food safety, environmental consequences, and regulations.

The Scale of the Global Food Waste Problem

Food waste remains one of the largest problems in the modern food system with inefficiencies being at the forefront Per Economic Research Service, USDA 133 billion pounds out of the estimated 430 billion pounds of food that was available at retail and consumer levels at the time was lost or discarded in year 2000 (31 percent of 133 billion pounds, which is 31% altogether, 133 + 430=563?). That equates to a retail value of $161.6 billion (J. Buzby, G.A. Farah-Wells, and M. Hyman). The United States Food and Drug Administration has cited the same number and, in addition, pointed out that food is not just the biggest landfill category nationally but the largest category altogether (USDA). EPA, in their most current study, indicates the numbers are still very big: Over a third of U.S. food goes to waste, and the average family of four loses nearly $3,000 in edible food annually. Deterioration and rotting of food are among the first reasons for loss. In particular, meat, fish, dairy, fruits, and vegetables have shorter shelf-lives.

Growing recognition that much of the waste happens while food is kept cold and transported to the market, rather than just overproduction, has made the development of packaging technology a point of great significance (J. Martin, C. W. Torres-Frenzel, and T. Wiedmann). On top of the financial loss caused by wasted food, the disposal of the resources that

went into producing and transporting it, i.e., water, land, energy, and labor, is a serious environmental issue. And, when food is dumped in the landfills and rotting there, methanol, which has a higher warming potential than carbon dioxide and is a potent greenhouse gas, is generated (USDA, 2003). Because of this, one might expect that reducing the amount of waste is beneficial to the environment as well as being helpful for the economy. Packaging that prevents spoilage and keeps the product in a desirable condition for a longer time. This will not only help individuals to manage their household budgets but will also affect the questions of resource saving and climate change (Rodrigues et al. 2017).

Why Conventional Packaging Has Limits

Packaging acts like a protective blanket by slowing factors like oxygen, moisture, and temperature that lead to the microbial decomposition of a product. This helps foods stay longer, and the number of throws away from these spoiled and/or non-consumable foods reduces as well. Spoilage factors that are most often related to spoilage are temperature, moisture, and microbiological growth; sometimes, food products may reach a level that is no longer safe or is not acceptable to customers even before showing signs of spoilage (Sooriyabandara and Kavirathna). Methods of traditional preservation, like vacuum sealing and modified atmosphere packaging, have brought an upgrade over more primitive ways, but these techniques mainly slow these changes rather than actively interfering with them (Jribi et al.). For very perishable foods, this difference is crucial: a packaging device that just limits the oxygen or moisture transfer will not have the same performance as the package that is actively holding back the growth of microorganisms or is giving a spoilage alarm in real-time. This performance gap is the reason, in part, that made a lot of nanotechnology solutions to food preservation highly promising from the food science and packaging manufacturing sides (Davidescu et al).

Antimicrobial Nanoparticles and Extended Shelf Life

The field of nanotechnology application related to food and its packaging, the research on the usage of nanoparticles of an antimicrobial nature (In particular silver, titanium dioxide, and zinc oxide) that could be introduced into the package is the widest known among others (Dash et al.). The microorganisms are inhibited when their cell membranes and metabolic processes are affected by these particulates, as they form reactive oxygen species and metals (Kotlhao et al.). Field reviews reveal that the addition of silver nanoparticles to packaging films has led to an increase from a mere 3-day shelf life to about 10 days for pork and comparable results in fish and dairy products that were subjected to refrigeration storage (Kraniewska Galus, Gniewosz). Some papers, on this have also detailed green-synthesized silver nanoparticle coating for starch-based packaging combining antibacterial properties toward E. coli plus a significant increase in mechanical strength (Rhim et al.) Whereas, research related to zinc oxide nanocomposites has reported antibacterial gains similar to those of silver nanoparticles with enhanced thermal stability in biodegradable films (Rhim Hong Ha). With a few more days of freshness, consumers

are not likely to throw away their product at either stores or households, so these antimicrobial functions lead to a significant drop in food waste (Amani, Gadde).

Improved Barrier Performance Through Nanocomposites

In addition to antimicrobial activity, nanocomposite materials - the hybrid structures formed as the combination of nanoparticles with plastics or biodegradable polymers - can enhance the structural aspect of packaging (Dash et al.). In such materials, the diffusion of oxygen as well as water vapour is hindered, which first slows the degradation of the food products and, on the other hand, makes the food packaging more durable and safe(Zhu et al.). Obviously, the use of biocomposite nanomaterials leads to a reduction of the rate of permeability of oxygen to the films and of the water vapour absorption, a property that has been found in combination with thermal stability versus the standard plastic films (Perera, Jaiswal, and Jaiswal). Nanocellulose, which is a biodegradable nanomaterial obtained from plant or bacterial cellulose, has recently been identified as a special kind of filler with many advantages for such nanocomposite materials (Thomas et al.). The reason is that it simultaneously offers great mechanical strength and a very effective water vapor-oxygen barrier; last, but not least, it is compostable (Hubbe et al.). Since these nanocomposite materials can be thinner and lighter yet perform as well, the amount of plastic needed for each package can be greatly reduced.

Intelligent and Smart Packaging Systems

A third application extends nanotechnology beyond passive protection into active monitoring. Nanosensors embedded in packaging can detect chemical changes associated with spoilage, such as the release of specific gases as bacteria break down food, and many of these systems produce a visible color change to alert consumers (Priyadarshi et al.). This kind of intelligent packaging addresses two problems simultaneously: it can reduce foodborne illness by flagging genuinely unsafe products, and it can reduce unnecessary waste caused by consumers discarding food that is still perfectly edible but past a printed date label (Hossain). Because confusing or overly conservative date labeling is itself a documented driver of household food waste, real-time spoilage indicators offer a more accurate and arguably more sustainable alternative to fixed expiration dates. (Hossain) The major applications of nanotechnology in food packaging and their relationship to food preservation and potential safety considerations are summarized in Figure 1.

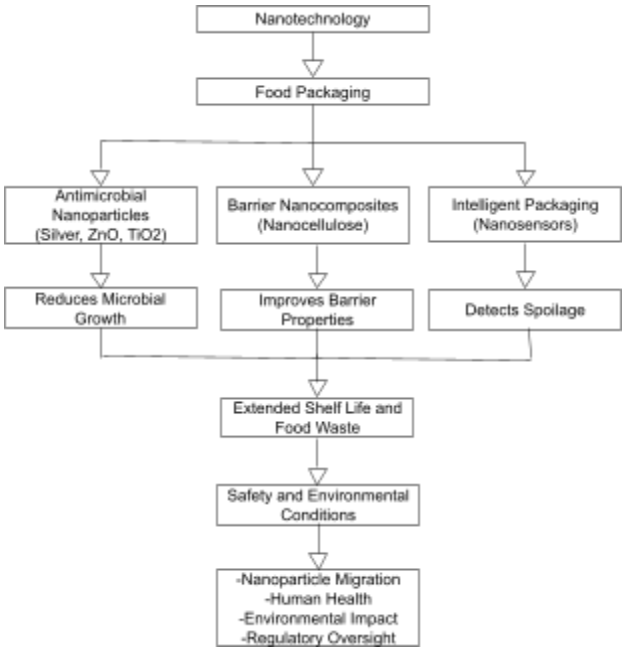


Figure 1. Overview of the major applications of nanotechnology in food packaging and their relationship to food preservation and potential safety considerations.

Table 1 compares the principal nanomaterials discussed in this review, categorizing them by their primary functions, advantages, and potential limitations.

Nanomaterial	**Primary Function**	**Major Benefits**	**Potential Risks**	**References**
Silver nanoparticles (AgNPs)	Antimicrobial packaging	Strong antimicrobial activity; extends shelf life	Potential migration into food; possible cytotoxicity following prolonged exposure	Kraśniewska et al. (2020); Rhim et al. (2010)
Zinc oxide (ZnO)	Antimicrobial packaging and UV protection	Inhibits bacterial growth; improves packaging durability	Possible oxidative stress; long-term safety requires further study	Kotlhao et al. (2017); Sharma et al. (2021)
Titanium dioxide (TiO_2)	UV protection and antimicrobial activity	Improves barrier properties; slows spoilage	Uncertainty regarding long-term ingestion and nanoparticle accumulation	Dash et al. (2022); Sharma et al. (2021)
Nanocellulose	Barrier nanocomposites	Biodegradable; excellent oxygen and moisture barrier; improved mechanical strength	Limited long-term commercial and toxicological data	Ahankari et al. (2021); Thomas et al. (2018)
Nanosensors	Intelligent food packaging	Detect spoilage; improve food safety monitoring; reduce unnecessary food waste	Higher production costs; limited large-scale commercial implementation	Priyadarshi et al. (2026); Khan et al. (2024)

Table 1. Comparison of major nanomaterials used in food packaging, their primary functions, benefits, and potential risks.

Safety and Migration Concerns

The same characteristics that turn nanotubes into powerful antimicrobiotics also make scientists worry about the possibility of these very tiny particles entering the human body (Ahankari et al.). Since the size of nanotubes is so small, they have the possibility of escaping from materials used for packaging food or into the food itself, for example, when food products are stored under high heat and acidity, or when the contact with the surface is of long duration (Sharma et al.). From the human point of view, after intake, some research shows that nanotubes can get across membranes more easily than bulk particles, thereby getting into the body and causing oxidative stress or inflammation that are observed in animal models after exposure to silver, titanium dioxide, or zinc oxide nanoparticles (Sarfraz et al.). A report focusing on migration mechanisms is also critical of the sensitivity of current methods in detection, meaning that they cannot accurately reveal the quantity of a particular nanoparticle that actually transfers, which makes it hard to establish clear boundaries of how safe a technology is (Ansari et al.). In addition, scientists have expressed concern about the gradual accumulation of low levels of nanotubes throughout life, mainly, of course, in children and pregnant women, as the effects of nanoparticle exposure in the aggregate on a single individual have not yet been studied (Sharma et al.). The uncertainties don't mean the nanoparticle-containing material itself is dangerous, but they do highlight the need for migration testing and toxicological work to be ahead of technological developments for which nano-based food-packaging materials would be used (Ahankari et al.).

Migration researchers further say that nanofillers are not only characterized by their particle size, but rather the amount (in percentage) of nanofiller in the composite material, the fat or water content level of food, temperature level, and the time the food stays in contact with the packaging are also factors that influence the extent of material migration (Akhter et al.). These are factors of variation that make it impossible for regulators to set a single worldwide safety limit. The result is that many toxicologists have insisted that nanotechnology-based foods or materials should each be assessed for safety purposes separately instead of one-time blanket approvals based on the type of particles (Omerovi et al).

Environmental Risks Beyond the Package

Concerns about nanoparticles go beyond just direct human contact. If nano-enhanced packaging is thrown away, its nanoparticles could end up in the landfills' leachate or sewages and that means they make their way to soils and waters (Keller et al). As we have limited knowledge on the fate of these particles once they are released to the environment (e.g., whether they remain, break down, or accumulate in organisms), scientists are advising that their impact on the ecosystem should be investigated more thoroughly before their extensive commercial usage (AZoNano). This environmental risk is a reflection of human health concerns, adding an extra level of seriousness to the uncertainty about how directly such particles affect humans (Ahankari et al).

Regulatory Oversight in the United States

Food packaging materials are evaluated by the Food and Drug Administration in the United States with a focus on two factors: first, the chance that a substance will get into food, and second, the amount of human exposure if it does (U.S. Food and Drug Administration). Those substances likely to transfer into food at a significant level would face additional safety requirements before getting permission. The agency has pointed out that materials might display different behavior at the nanoscale than in their bulk form, so requiring additional testing (U.S. Food and Drug Administration). When it comes to nanomaterials in food, though some critics maintain that U. S. regulation basically depends on voluntary industry consultation, without enforcing binding pre-market requirements, thereby creating loopholes in the degree of consistency with which nanoscale ingredients and packaging components are actually checked for safety before they get onto consumers' tables (U.S. Food and Drug Administration).

Regulatory Oversight in the European Union

On the one hand, the EU has taken more systematic regulatory routes to ensure consumers' safety by setting strict controls on food contact materials. Food contact materials in the EU are regulated under a set of comprehensive regulations laying out the maximum level to which a substance may migrate into food, with special provisions on active and intelligent packaging technologies that could affect the migration (European Commission). Besides that, the European Food Safety Authority (EFSA) has drafted an informative text that describes how nanomaterials have to be characterized and the risk assessment to be done. In particular, it lists the main conditions when it becomes necessary to test nanospecific properties of a material's small particle fraction, beyond a conventional safety assessment (EFSA Scientific Committee). Such risk analysis would apply to various EFSA areas, including novel foods, food additives, and food contact materials. In this way, they intend to implement a uniform approach to risk-based decision-making for any nanomaterial, no matter through which sector it will enter the food chain. (EFSA Scientific Committee).

The Challenge of Global Regulatory Harmonization

As the United States, the European Union, and other parts of the world have crafted varied regulations on what is considered a nanomaterial and what kind of tests products should go through, the packaging sector's nano-enabled products could find it quite difficult to market internationally because of an extremely complex and disjunct regulatory setup (Anand et al.). Besides, differences in standards might obstruct the deployment of promising new technologies, lead to a higher compliance burden for the companies, and pose difficulties in achieving the same level of consumer safety across different markets (OECD). Several scientists working in the area are convinced that common testing procedures, definitions of nanomaterials agreed upon by all parties involved in their design and production, and a more efficient exchange of regulatory information between authorities would be necessary if we are to combine our ambition to foster technological progress with our safety concerns (Allan et al.).

Future Research Needs and Perspectives

Nanocellulose, among the nanomaterials presented in this review, is one of the safety profile-friendly materials because it is degradable, naturally sourced, and it prevents the escape of gases effectively without the dependency on metal-based antimicrobials (Ahankart et al.). By comparison, silver nanoparticles are the most effective but also the safest, as they present problems related to the movement of nanoparticles and toxicity over time. Zinc oxide and titanium dioxide could be considered to strike the best compromise of the two, with antimicrobial effectiveness and safety level. Still, more long-term investigations will have to confirm this (Rao et al.). Among the applications currently under development, antimicrobial packaging and nanocomposite barrier material look most ready for commercial applications, while intelligent nanosensor systems remain less popular owing to high production costs and technological complexity (Khan et al.).

Significant research work has been done; Yet, there are still many gaps. Future research should address issues such as long-term toxic effects, developing standard ways of measuring nanoparticle migration, and assessing the environmental fate of disposed of nano-enabled packaging materials (Gupta et al.). To develop harmonized regulatory standards that ensure protection of public health as well as promote the development of new types of food packaging technologies, more international cooperation is necessary.

Conclusion

The development of the potential offered by nanotechnology has brought about the idea that nanotechnology is a genuine way to solve the food packaging problems related to shelf life, spoilage, and monitoring food safety. Food waste is reduced directly, a problem that affects the US economy very badly in the sense that the loss exceeds $ 100 billion each year. In parallel, the migration of the nanoparticles into the food and the environment is still not well known, and the current detection and testing methods have not kept up with the pace of commercial usage. It means that the technology is not yet a breakthrough nor a hazard; it is a new field, and what will eventually change it is the toxicological research that will provide the quality data, the regulations that will be strict in the US and EU, and the other regions, and the cooperation of the industrial and governing bodies in the public communication. Nanotechnology should, given the evidence available at present, proceed along its developmental path with a minimum of disturbance until the long-term toxicological studies, the standardized migration testing, and the environmental assessments have given their results. The priority of future research must be to increase the commercial viability of the smart nanosensor technologies and to enhance international collaboration for setting up the internationally harmonized regulatory standards that are a combination of protecting public health and stimulating continuous innovation in food packaging. Setting the international rules and standards of regulation, and the expansion of the safety evaluations based on evidence, will be crucial in making sure that the gains from food packaging enhanced by nanotechnology can be achieved without harming human health or the environment.

Works Cited

Buzby, Jean C., Hodan Farah-Wells, and Jeffrey Hyman. “The Estimated Amount, Value, and Calories of Postharvest Food Losses at the Retail and Consumer Levels in the United States. SSRN, 2014,” Accessed April 29, 2026.https://doi.org/10.2139/ssrn.2501659.

United States Department of Agriculture. “Why Should We Care About Food Waste?” U.S. Department of Agriculture. Accessed April 29, 2026. https://www.usda.gov/about-food/food-safety/food-loss-and-waste/why-should-we-care-about-food-waste.

Kenny, Shannon. Estimating the Cost of Food Waste to American Consumers. U.S. Environmental Protection Agency, Office of Research and Development, EPA/600/R-25/048, Apr. 2025. Accessed April 29, 2026. https://www.epa.gov/system/files/documents/2025-04/costoffoodwastereport_508.pdf.

Martin, Nicole H., Paula Torres-Frenzel, and Martin Wiedmann. “Invited Review: Controlling Dairy Product Spoilage to Reduce Food Loss and Waste.” Journal of Dairy Science, vol. 104, no. 2, 2021, pp. 1251–126. Accessed April 30, 2026. https://doi.org/10.3168/jds.2020-19130.

United States Department of Agriculture. “Food Loss and Waste.” U.S. Department of Agriculture. Accessed April 30, 2026.https://www.usda.gov/foodlossandwaste/consumers.

Rodrigues, Vanessa, Eileen Cahill, and Karina E. Mis Solval. “The Potential for Reducing Food Waste through Shelf-Life Extension: Actionable Insights from Data Digitization.” Sustainability, vol. 16, no. 7, 2024, article 2986. Accessed May 2, 2026. https://doi.org/10.3390/su16072986.

Sooriyabandara, W. T. N., and C. A. Kavirathna. “The Role of Logistics in Perishable Food Waste Reduction of Small and Medium-Sized Enterprises: A Systematic Review.” 2024 International Conference on Advances in Technology and Computing (ICATC), IEEE, 2024, pp. 1–6. Accessed May 2, 2026. https://doi.org/10.1109/ICATC64549.2024.11025386.

Jribi, Sarra, et al. “COVID-19 Virus Outbreak Lockdown: What Impacts on Household Food Wastage?” Environment, Development and Sustainability, vol. 22, no. 5, 2020, pp. 3939–3955. Springer. Accessed May 2, 2026.https://doi.org/10.1007/s10668-020-00740-y.

Davidescu, Mădălina Alexandra, et al. “Advances and Challenges in Smart Packaging Technologies for the Food Industry: Trends, Applications, and Sustainability Considerations.” Foods, vol. 14, no. 24, 2025, article 4347, MDPI. Accessed May 2, 2026. https://doi.org/10.3390/foods14244347.

Dash, Kshirod Kumar, et al. “Applications of Inorganic Nanoparticles in Food Packaging: A Comprehensive Review.” Polymers, vol. 14, no. 3, 2022, article 521. MDPI. Accessed May 3, 2026. https://doi.org/10.3390/polym14030521.

Kotlhao, Kate, et al. “The Synthesis of Silver, Zinc Oxide and Titanium Dioxide Nanoparticles and Their Antimicrobial Activity.” Advanced Materials Proceedings, vol. 2, no. 8, 2017, pp. 479–484. VBRI Press. Accessed May 4, 2026.https://doi.org/10.5185/amp.2017/803.

Kraśniewska, Karolina, Sylwia Galus, and Małgorzata Gniewosz. “Biopolymers-Based Materials Containing Silver Nanoparticles as Active Packaging for Food Applications—A Review.” International Journal of Molecular Sciences, vol. 21, no. 3, 2020, article 698. MDPI. Accessed May 4, 2026. https://doi.org/10.3390/ijms21030698.

Rhim, Jong-Whan, et al. “Silver Nanoparticle-Loaded Chitosan–Starch Based Films: Fabrication and Evaluation of Tensile, Barrier and Antimicrobial Properties.” Materials Science and Engineering: C, vol. 30, no. 6, 2010, pp. 891–897. Elsevier. Accessed May 5, 2026. https://doi.org/10.1016/j.msec.2010.04.004.

Rhim, Jong-Whan, Sang-Ho Hong, and Ha-Jung Ha. “Bio-Nanocomposites for Food Packaging Applications.” Carbohydrate Polymers, vol. 90, no. 1, 2012, pp. 467–473. Elsevier. Accessed May 15, 2026. https://doi.org/10.1016/j.carbpol.2013.01.007.

Amani, Pegah, and Lars-Erik Gadde. “Shelf Life Extension and Food Waste Reduction.” EFITA-WCCA-CIGR Conference Proceedings, edited by Gerhard Schiefer and Ursula Rickert, 2015, pp. 7–14. AgEcon Search. Accessed May 15, 2026. https://doi.org/10.22004/ag.econ.206209.

Dash, Kshirod Kumar, et al. “Applications of Inorganic Nanoparticles in Food Packaging: A Comprehensive Review.” Polymers, vol. 14, no. 3, 2022, article 521. MDPI. Accessed May 16, 2026.https://doi.org/10.3390/polym14030521.

Zhu, Junyong, Nicole M. Stark, Peter Kitin, Nayomi Z. Plaza, and Jenna Zukswert. “Cellulose-Based Barrier Coatings for Fruit: Using Hardwood Pulp to Reduce Food Waste.” Rooted in Research, no. 56, U.S. Department of Agriculture, Forest Service, Forest Products Laboratory, May 2026. Accessed May 16, 2026. https://research.fs.usda.gov/treesearch/80566.

Perera, Kalpani Y., Amit K. Jaiswal, and Swarna Jaiswal. “Biopolymer-Based Sustainable Food Packaging Materials: Challenges, Solutions, and Applications.” Foods, vol. 12, no. 12, 2023, article 2422. MDPI. Accessed May 16, 2026. https://doi.org/10.3390/foods12122422.

Thomas, Bejoy, et al. “Nanocellulose, a Versatile Green Platform: From Biosources to Materials and Their Applications.” Chemical Reviews, vol. 118, no. 24, 2018, pp. 11575–11625. American Chemical Society. Accessed May 19, 2026. https://doi.org/10.1021/acs.chemrev.7b00627.

Sarfraz, Jawad, et al. “Nanocomposites for Food Packaging Applications: An Overview.” Nanomaterials, vol. 11, no. 1, 2021, article 10. MDPI. Accessed May 19, 2026. https://doi.org/10.3390/nano11010010.

Priyadarshi, Ruchir, et al. “Toxicity and Safety of Nanomaterials in Food Packaging Applications: A Dimension-Based Review.” Comprehensive Reviews in Food Science and Food Safety, vol. 25, no. 1, 2026, e70374. Wiley. Accessed May 19, 2026. https://doi.org/10.1111/1541-4337.70374.

Hossain, Abul. “Do Date Codes Cause Food Waste? Smart Packaging Might Tackle the Problem.” Journal of Food Bioactives, 2023, pp. 6–12. International Society for Nutraceuticals & Functional Foods. Accessed May 28, 2026.

https://doi.org/10.31665/JFB.2023.18359.
Ahankari, Sandeep S., Aditya R. Subhedar, Swarnim S. Bhadauria, and Alain Dufresne. "Nanocellulose in Food Packaging: A Review." Carbohydrate Polymers, vol. 255, 2021, article 117479. Elsevier. Accessed May 28, 2026. https://doi.org/10.1016/j.carbpol.2020.117479.
Sharma, Annu, et al. "Nanotechnology: An Untapped Resource for Food Packaging." Foods, vol. 10, no. 9, 2021, article 2114. MDPI. Accessed May 28, 2026. https://doi.org/10.3390/foods10092114.
Sarfraz, Jawad, et al. "Nanocomposites for Food Packaging Applications: An Overview." Nanomaterials, vol. 11, no. 1, 2021, article 10. MDPI. Accessed May 30, 2026. https://doi.org/10.3390/nano11010010.
Ansari, Mohammed Dilsad Izrayeel, Laura Principato, Luigi De Nardo, and Carlo Punta. "Nanomaterials in Food Packaging: An Overview of Regulatory Frameworks and Migration Assessment." Food Control, vol. 181, 2026, article 111707. Elsevier. Accessed May 30, 2026. https://doi.org/10.1016/j.foodcont.2025.111707.
Akhter, Sumi, et al. "Migration of Chemical Compounds from Packaging Materials into Packaged Foods: Interaction, Mechanism, Assessment, and Regulations." Foods, vol. 13, no. 19, 2024, article 3125. MDPI. Accessed May 30, 2026. https://doi.org/10.3390/foods13193125.
Omerović, Nermina, et al. "Nanomaterials in Food Packaging: Assessment of Environmental and Health Risks." International Journal of Environmental Research and Public Health, vol. 18, no. 13, 2021, article 6663. MDPI. Accessed June 5, 2026. https://doi.org/10.3390/ijerph18136663.
Keller, Arturo A., Suzanne McFerran, Anastasiya Lazareva, and Sangwon Suh. "Global Life Cycle Releases of Engineered Nanomaterials." Journal of Nanoparticle Research, vol. 15, no. 6, 2013, article 1692. Springer. Accessed June 5, 2026. https://doi.org/10.1007/s11051-013-1692-4.
AZoNano. "Nanotechnology in Food Packaging." AZoNano, 13 Nov. 2019. Accessed June 5, 2026.https://www.azonano.com/article.aspx?ArticleID=5114.
U.S. Food and Drug Administration. "Food Packaging & Other Substances that Come in Contact with Food – Information for Consumers." U.S. Food and Drug Administration. Accessed June 6, 2026. https://www.fda.gov/food/food-ingredients-packaging/food-packaging-other-substances-come-contact-food-information-consumers.
U.S. Food and Drug Administration. Guidance for Industry: Preparation of Food Contact Substance Notifications (Toxicology Recommendations). U.S. Department of Health and Human Services, Oct. 2021. Accessed June 10, 2026. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/guidance-industry-preparation-food-contact-substance-notifications-toxicology-recommendations.

U.S. Food and Drug Administration. “FDA’s Approach to Regulation of Nanotechnology Products.” U.S. Food and Drug Administration. Accessed June 10, 2026. https://www.fda.gov/science-research/nanotechnology/fdas-approach-regulation-nanotechnology-products.

European Commission. “Legislation.” Food Safety, European Commission. Accessed June 12, 2026.https://food.ec.europa.eu/food-safety/chemical-safety/food-contact-materials/legislation_en.

EFSA Scientific Committee. “Guidance on Risk Assessment of Nanomaterials to Be Applied in the Food and Feed Chain: Human and Animal Health.” EFSA Journal, vol. 19, no. 8, 2021, article e06768. European Food Safety Authority. Accessed June 12, 2026. https://doi.org/10.2903/j.efsa.2021.6768.

Anand, Avinash, Harmeet Kaur, Hunmily Hansepi, and Kaustubh Chandrakant Khaire. “Green Nanotechnology-Enabled Biopolymer Nanocomposites for Eco-Innovative and Regulatory-Compliant Antimicrobial Food Packaging.” Frontiers in Industrial Microbiology, vol. 4, 2026, article 1801909. Frontiers Media. Accessed June 21, 2026. https://doi.org/10.3389/finmi.2026.1801909.

Organisation for Economic Co-operation and Development. Recommendation of the Council on the Safety Testing and Assessment of Manufactured Nanomaterials. OECD Legal Instruments, OECD/LEGAL/0400, Organisation for Economic Co-operation and Development, 2025. Accessed June 21, 2026. https://legalinstruments.oecd.org/public/doc/298/298.en.pdf.

Allan, Jacqueline, et al. “Regulatory Landscape of Nanotechnology and Nanoplastics from a Global Perspective.” Regulatory Toxicology and Pharmacology, vol. 122, 2021, article 104885. Elsevier. Accessed June 22 2026.https://doi.org/10.1016/j.yrtph.2021.104885.

Rao, Madhushree M. V., Naeem Mohammad, Shaibal Banerjee, and Pawan K. Khanna. “Synthesis and Food Packaging Application of Silver Nano-Particles: A Review.” Sustainable Food Technology, vol. 2, 2024, pp. 1153–1176. Elsevier. Accessed July 12, 2026.https://doi.org/10.1016/j.sfot.2024.100205.

Khan, Shadman, et al. “Smart Food Packaging Commercialization.” Nature Reviews Bioengineering, vol. 2, 2024, pp. 535–537. Springer Nature. Accessed July 12, 2026. https://doi.org/10.1038/s44222-024-00190-5.

Gupta, Rakesh Kumar, et al. “Migration of Chemical Compounds from Packaging Materials into Packaged Foods: Interaction, Mechanism, Assessment, and Regulations.” Foods, vol. 13, no. 19, 2024, article 3125. MDPI. Accessed July 12, 2026. https://doi.org/10.3390/foods13193125.

A Bioinformatics-Guided Study of Moringa–Plant Polymer Coagulants for Sustainable Microplastic Removal from Water

By Neil Lewis[1] and Nirupma Singh[2]

Abstract

Microplastics have become an important environmental issue in water resources globally, especially in low-resource and community settings where advanced water treatment infrastructure is unavailable and current treatment technologies have shown limited ability to effectively remove microplastic particles. In this study, the potential of natural coagulant-based fabric filters for microplastic removal was computationally explored by probing the binding interactions between the *Moringa oleifera* Coagulant Protein (MOCP) and four microplastic surrogates namely polyethylene (PE), polyethylene terephthalate (PET), polypropylene (PP) and polystyrene (PS), using AutoDock Vina, and between four plant polymers i.e. pectin, starch, alginate and cellulose and the PET surrogate terephthalic acid using SMINA, both on the Tamarind Bio platform, with three-dimensional interactions visualised using PyMOL. Among the microplastic surrogates docked to MOCP, PET surrogate terephthalic acid demonstrated the greatest binding affinity of -5.019 kcal/mol, followed by PS surrogate styrene with -4.637 kcal/mol. Cellulose had the highest binding affinity for the PET surrogate terephthalic acid (-2.2 kcal/mol) among the plant polymers, followed closely by starch (-2.1 kcal/mol). These results suggest that *Moringa oleifera* Coagulant Protein and cellulose derivatives have the potential to coagulate microplastics, providing a promising computational basis for the development of biodegradable, low-cost fabric filters suitable for community and household water treatment systems.

Keywords *Moringa oleifera*, Microplastics, Plant polymers, Bioflocculation, Natural coagulants, Molecular docking, Cheminformatics, Water remediation

Introduction

Over the past few decades, global plastic production has drastically increased, and due to poor waste management, these particles have spread across oceans, rivers, soils, and even drinking water sources worldwide. Microplastics are plastic particles that are smaller than 5 mm in size and are produced either from the breakdown of larger plastic waste or from products like cosmetic microbeads. Microplastics were detected in tap water, bottled water, and freshwater bodies across multiple countries, raising serious concerns about human health. Once ingested, they were found to cause inflammation, hormonal disruption, and damage to multiple organs. In India, the problem was found to be particularly serious due to high plastic use and low recycling rates, with microplastic contamination being reported in major rivers, lakes, and even bottled drinking water brands [1, 2].

[1]Jamnabai Narsee International School, Mumbai, Maharashtra, India
[2] Department of Biological Sciences and Engineering, Netaji Subhas University of Technology (Formerly NSIT), Dwarka, Delhi, India.

Microplastic removal through coagulation using *Moringa oleifera* seed extract and anionic polyacrylamide was studied by comparing three polymer types, mainly polyamide, polystyrene, and polyethylene. The combination of aluminium sulphate with Moringa extract improved overall removal efficiency compared to either agent used alone, proving Moringa as an effective natural coagulant for microplastic treatment.[3] A literature review on wastewater treatment plants in Southeast Asia looked at how natural coagulants like *Moringa oleifera* and banana peel extracts remove microplastics. It was found that these natural options are better than chemical ones because they are non-toxic and cheaper.[4] A review on removing microplastics from different types of water studied inorganic, organic, and natural coagulants and found that natural ones like chitosan, starch, and protein amyloid fibrils worked well, but still need improvement for very small or irregularly shaped microplastics.[5] Water purification using *Moringa oleifera* seed paste combined with scoria filtration media was tested as a point-of-use treatment system. The combined coagulation-filtration approach significantly reduced turbidity and biological contamination, supporting the practical use of Moringa seed as a natural coagulant in low-resource community settings. [6] This review examined how coagulation is used to remove microplastics in water treatment plants, explaining charge neutralisation, adsorption bridging, and sweep flocculation. It also found that factors such as the type of coagulant, the size and shape of microplastics, pH, and dosage all influence the effectiveness of the removal. [7] Cationic-modified starch was developed and tested as a biodegradable coagulant to remove polystyrene microplastics from water and achieved about 65% removal, worked better for larger and aged particles, and was effective across a wide pH range, making it a promising eco-friendly alternative to chemical coagulants.[8]

The coagulation of pristine and weathered polyethylene, polystyrene, and polyester microplastics was studied using jar tests and quartz crystal microbalance. It was also found that cationic polyacrylamide formed flocs more effectively than aluminium-based coagulants.[9] A study used protein amyloid fibrils as a natural coagulant to remove microplastics and organic matter from water. The fibrils performed better than traditional metal-based coagulants, achieving high removal efficiency even at low doses. It showed that protein-based natural coagulants are a promising and sustainable option for water treatment.[10] A biochar filter made from banana peel was tested as a continuous-flow system to remove polystyrene microplastics. Factors like bed height, flow rate, and particle size were studied, and over 92% removal was achieved under optimal conditions. This shows that low-cost materials like agricultural waste can be effective for microplastic filtration.[11] A combined coagulation system using ferric chloride and chitosan was tested to remove polystyrene microplastics from wastewater. It achieved up to 99.8% removal under optimal conditions, performing much better than either one alone. Artificial neural networks were also used to predict the best treatment conditions.[12] Polyvinyl chloride microplastics were removed from drinking water using *Moringa oleifera* extract and aluminium sulphate. Both achieved over 98% removal, with Moringa performing as well as or better than alum in alkaline conditions, showing it as an effective plant-based coagulant.[13] A new coagulation method using metal-phenolic bonds was tested for microplastic removal. Unlike

traditional iron or aluminium salts, tannic acid–based coagulation achieved much higher removal across different microplastics. Toxicity tests also showed the treated water was safe. [14]

Many research gaps were found across the reviewed literature, almost all studies were done in labs using synthetic water, natural coagulants were only tested as loose additions to water and never embedded into a fabric, most of the studies focused on just one or two types of plastic, and very few considered if the system was affordable or biodegradable enough for real-world use, especially in rural areas of India. To address these gaps, it was hypothesised that fabric filters embedded with natural coagulants like *Moringa oleifera* seed extract or starch derivatives would cause small plastic particles to clump together and be captured within the fabric, effectively removing microplastics from water. It was further expected that this could be tested and confirmed using light microscopy or simple imaging tools, and that the final product would work as a biodegradable, low-cost filter that could easily be integrated into home or community water systems.

Methodology

Data Collection

An experiment on the use of *Moringa oleifera* seed extract and scoria rock filter was applied in a point-of-use water treatment experiment where the removal of turbidity, phosphate, fluoride, and iron were analyzed depending on the different doses and time of exposure. A machine learning experiment based on 527 cases of water treatment plant data was conducted in comparing the efficiency of flocculation when *Moringa oleifera* is compared with chitosan in varied doses, pH, and turbidity. Another experiment using Moringa peregrina for the purpose of analyzing its efficiency as a bio-coagulant was conducted to compare its effectiveness against raw seeds, raw extract, and de-oiled extract. The effectiveness of nopal mucilage and cassava starch was compared with *Moringa oleifera* in removing turbidity using experiments as a natural coagulant based on starch. The analysis of microplastics in drinking water was carried out based on data derived from a dataset including data from 46 research works which included both tap water and bottled water, along with pre-packaged beverage cartons. The data used for studying the distribution of microplastics in oceans included a dataset with 8,218 samples of microplastics gathered across the globe from oceans from the years 2000 till 2019 and another database with more than 22,500 records regarding marine microplastics. Global riverine plastic pollution data was taken into account to determine the level of plastic pollution of water bodies. For determining water quality of Indian rivers and water bodies, data regarding pH, dissolved oxygen, conductivity, BOD, nitrate content, and fecal coliforms was obtained from government websites. The drinking water filtration experiment done on harvested rainwater was taken into consideration to evaluate the quality of water at three phases –first flush, pre-filtered, and filtered.

Molecular Docking

Molecular docking was conducted to analyze the interaction between the protein and the ligand. The protein structure was sourced from the RCSB PDB (Protein Data Bank) which is an online database for storing structural information of biological macromolecules while the ligands were sourced from PubChem, which is an online database of NCBI(National Center for Biotechnology Information) containing information on chemical structure and properties of small chemicals. The *Moringa oleifera* Coagulant Protein (MOCP) receptor was uploaded to the Tamarind Bio AutoDock Vina software platform, and four ligands including Polyethylene (PE) surrogate butane (CID 7843), Polyethylene Terephthalate (PET) surrogate terephthalic acid (CID 7489), Polypropylene (PP) surrogate propylene (CID 8252), and Polystyrene (PS) surrogate styrene (CID 7501) were separately docked with the protein. The binding pocket was determined manually by moving the X, Y, and Z coordinates of the docking box such that it surrounds the protein active site. Then each of the ligands was docked separately against the protein. The greater the negativity of the score, the greater the binding affinity between protein and ligand.

Another molecular docking experiment was performed to determine the interaction of the plant polymer and the microplastic surrogate. Microplastics surrogates, namely Polyethylene (PE) surrogate butane (CID 7843), Polyethylene Terephthalate (PET) surrogate terephthalic acid (CID 7489), Polypropylene (PP) surrogate propylene (CID 8252), and Polystyrene (PS) surrogate styrene (CID 7501), which are available from the online PubChem database from the National Center for Biotechnology Information (NCBI) that gives information on the chemical structure and property of small molecules, were used as the receptor. Similarly, Pectin (CID 441476), Alginate (CID 91666324), Starch (CID 47207535), and Cellulose (CID 16211032), which are available from PubChem database were utilized as the ligand. Individual docking of each microplastic surrogate to the four plant polymers as the ligand was done separately using the SMINA docking tool on the Tamarind Bio platform giving rise to a total of 16 docking interactions. Binding pocket determination was done manually by manipulating the X, Y, and Z-coordinates of the docking box surrounding the active site of the microplastic surrogate. Binding affinity values in kcal/mol were recorded where the lower the value, the greater the binding affinity.

PyMol visualisation

The docked results obtained from Tamarind Bio were imported into PyMOL, which is a molecular visualization software tool to visualize the three-dimensional interaction of protein and ligand molecules. The docked receptor file and the file containing the best ligand pose were loaded into PyMOL to be able to visualize the protein-ligand complexes of MOCP and microplastics in three-dimensional space. A similar step was followed for the plant polymer and microplastic complexes. The distances between the ligand atoms and the receptor present in the binding site were calculated using the Wizard tool. The Wizard tool is an inbuilt measuring tool in PyMOL, which measures the distance between two selected atoms in Ångströms (Å). A distance of 5 Å was set as the criterion for meaningful interactions, because any distance below this is an indication of close contact and molecular interaction.

Results

Molecular Docking Analysis of MOCP with Microplastics

Molecular docking was done to study the interaction of the *Moringa oleifera* Coagulant Protein (MOCP) and four microplastic surrogates in an attempt to determine computationally whether the protein could bind to microplastics of different compositions. Molecular docking is a computational process which attempts to find out how two molecules can fit together in three dimensional space and how much affinity there will be between them. In the current study, the protein receptor, 6S3F Moringa MOCP, was docked to four ligands; polyethylene (PE) surrogate butane (CID 7843), polyethylene terephthalate (PET) surrogate terephthalic acid (CID 7489), polypropylene (PP) surrogate propylene (CID 8252) and polystyrene (PS) surrogate styrene (CID 7501) using the AutoDock Vina program on Tamarind Bio. A docking box was defined within the binding pocket of the protein by setting the X, Y, and Z axis coordinates to 102, 12, and 11 respectively, with a width of 37, height of 48, and depth of 33. The binding affinities were -2.621 kcal/mol for PE surrogate butane, -5.019 kcal/mol for PET surrogate terephthalic acid, -2.015 kcal/mol for PP surrogate propylene and -4.637 kcal/mol for PS surrogate styrene, as shown in Table 1. Of all the ligands, the PET surrogate terephthalic acid showed the highest binding affinity to the Moringa MOCP receptor followed by PS surrogate styrene.

Microplastic type	**Surrogate ligand**	**CID**	**Receptor**	**Binding affinity (kcal/mol)**
Polyethylene (PE)	Butane	7843	6S3F MOCP	-2.621
Polyethylene Terephthalate (PET)	Terephthalic acid	7489	6S3F MOCP	-5.019
Polypropylene (PP)	Propylene	8252	6S3F MOCP	-2.015
Polystyrene (PS)	Styrene	7501	6S3F MOCP	-4.637

Table 1. Binding Affinities of MOCP with Microplastics using molecular docking.

Molecular Docking Analysis of Plant Polymers with Microplastics

Molecular docking was done to examine the interaction of four plant polymers with four microplastic surrogates to determine the possibility that plant-based polymers could interact with different types of microplastics through the docking process computationally. In the present study, the four microplastic surrogates such as polyethylene (PE) surrogate butane (CID 7843), polyethylene terephthalate (PET) surrogate terephthalic acid (CID 7489), polypropylene (PP) surrogate propylene (CID 8252), and polystyrene (PS) surrogate styrene (CID 7501) were used as the receptors, whereas pectin (CID 441476), alginate (CID 91666324), starch (CID 47207535), and cellulose (CID 16211032) were used as ligands. After performing the docking of each individual microplastic surrogate with all four plant polymer ligands through the SMINA docking tool on the Tamarind Bio platform, it generated a total of 16 docking complexes. All 16

binding affinity scores obtained from the docking results were reported in kcal/mol, where a more negative score indicates stronger binding affinity. The molecular docking displayed that PET, out of all the microplastics, had the highest binding affinity with all the plant polymers i.e. -1.8 kcal/mol with Pectin, -2.1 kcal/mol with Starch, -1.7 kcal/mol with Aliginate, and -2.2 kcal/mol with Cellulose. All the 16 binding affinities of the plant polymers with microplastics are shown in Table 2.

Plant Polymer (Ligand)	**Microplastic Surrogate (Receptor)**	**Binding Affinity (kcal/mol)**
Pectin	PE surrogate butane	-0.6
Pectin	PET surrogate terephthalic acid	-1.8
Pectin	PP surrogate propylene	-0.5
Pectin	PS surrogate styrene	0
Starch	PE surrogate butane	-1.0
Starch	PET surrogate terephthalic acid	-2.1
Starch	PP surrogate propylene	-0.9
Starch	PS surrogate styrene	-1.9
Alginate	PE surrogate butane	-0.6
Alginate	PET surrogate terephthalic acid	-1.7
Alginate	PP surrogate propylene	-0.5
Alginate	PS surrogate styrene	-1.1
Cellulose	PE surrogate butane	-0.9
Cellulose	PET surrogate terephthalic acid	-2.2
Cellulose	PP surrogate propylene	-0.8
Cellulose	PS surrogate styrene	-1.6

Table 2. The Binding Affinities of all the plant polymers with microplastics.

Interaction Analysis in PyMOL

PyMOL is a software for molecular visualization and was employed to investigate the 3D interactions between the docked protein-ligand complexes. This program is especially helpful as it provides the opportunity to look at the spatial organization of the molecules and measure distances between interacting atoms. The docked output files from Tamarind Bio were opened in

PyMOL and the receptor and ligand poses were chosen for visualization. In order to measure the distances between interacting atoms in the binding site, the Wizard tool in PyMOL was used, which is an inbuilt measurement feature that measures the distance between two selected atoms in Ångströms (Å). Three distances were measured for each complex, and the threshold value was set to 5 Å, as any distance below this value was considered an indicator of close contact and interaction between the two molecules. The measured distances are shown by yellow dotted lines with their values indicated on the structure. For the MOCP receptor docked with the four microplastic surrogates, the visualisations are shown in Figure 1 (a–d). For the plant polymer ligands namely pectin, starch, alginate, and cellulose, docked with the PET surrogate terephthalic acid as the receptor, the visualisations are shown in Figure 2 (a–d).

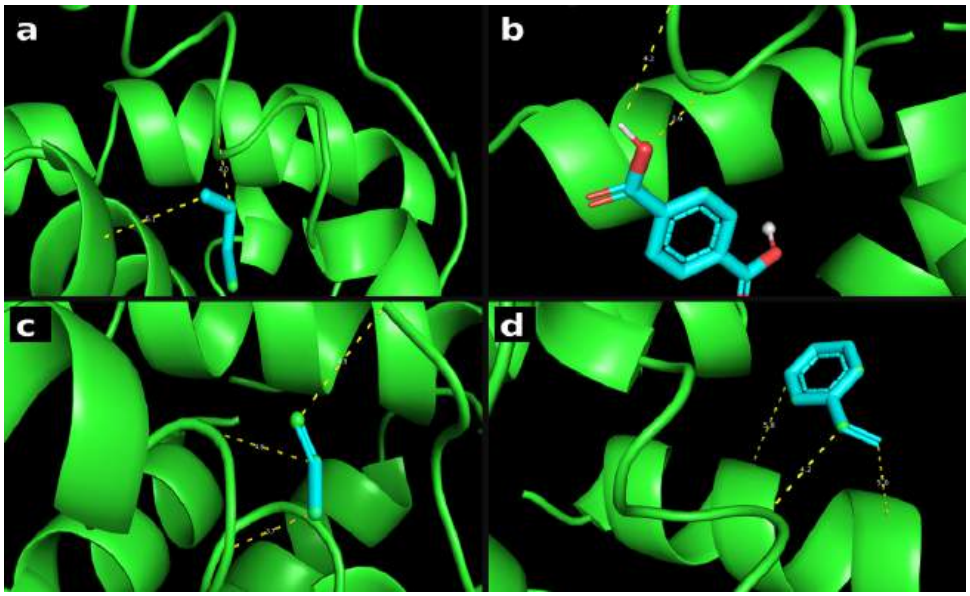


Figure 1. PyMOL visualisation of *Moringa oleifera* Coagulant Protein (MOCP) receptor docked with four microplastic surrogates. (a) MOCP docked with PE surrogate butane, (b) MOCP docked with PET surrogate terephthalic acid, (c) MOCP docked with PP surrogate propylene, (d) MOCP docked with PS surrogate styrene. The green ribbon structure represents the MOCP receptor and the cyan stick structure represents the docked ligand. Yellow dashed lines indicate measured distances in Ångströms (Å) between the ligand and nearby protein residues at the binding site.

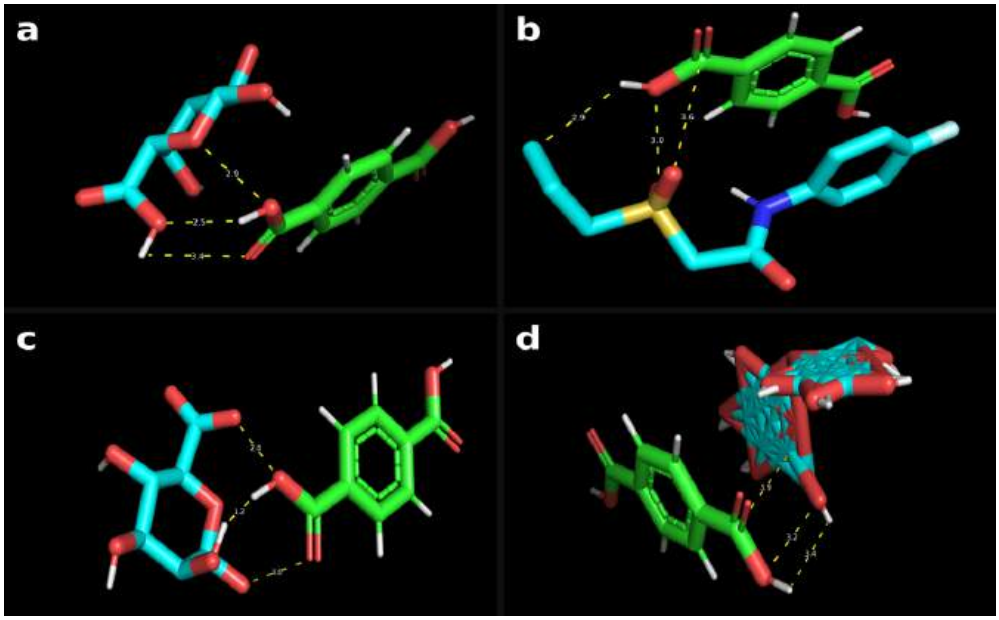


Figure 2. PyMOL visualisation of plant polymer ligands docked with microplastic surrogates. (a) Pectin docked with microplastic surrogate, (b) Starch docked with microplastic surrogate, (c) Alginate docked with microplastic surrogate, (d) Cellulose docked with microplastic surrogate. The cyan stick structure represents the plant polymer ligand and the green stick structure represents the microplastic surrogate receptor. Yellow dashed lines indicate measured distances in Ångströms (Å) between interacting atoms at the binding site.

Discussion

For the examination of the potential interaction of the *Moringa oleifera* Coagulant Protein (MOCP) with four microplastic surrogates and for the four plant polymers with four microplastic surrogates, molecular docking was used as a computational approach in the present experiment. The MOCP receptor was selected because this coagulant protein is the active compound of *Moringa oleifera* seeds which has water purification ability. Instead of entire polymer chains, microplastic surrogates were utilized, which can represent the chemical structure of each plastic type in a computationally manageable format. Among the four microplastics docked with MOCP, PET surrogate terephthalic acid had the strongest binding affinity at −5.019 kcal/mol, followed by PS surrogate styrene at −4.637 kcal/mol, suggesting that MOCP tends to interact more strongly with aromatic and polar microplastic types. Of the four plant polymers docked against all four microplastic surrogates, cellulose and starch had the highest binding affinities with the PET surrogate terephthalic acid at -2.2 kcal/mol and -2.1 kcal/mol, respectively, implying these two plant-based polymers have the greatest potential to bind with PET microplastics.

In a previous study that performed molecular docking of *Moringa oleifera* compounds on the apoptosis protein p53, AutoDock was used to determine the binding affinities of Moringa compounds to a protein from the RCSB-PDB, as in the present study. However, that study carried out molecular docking to determine the anticancer activity of Moringa compounds, making the research problems fundamentally different, as although the same method and plant

were used, the application was completely different.[15] Another study that performed molecular docking of *Moringa oleifera* compounds against SARS-CoV-2 proteins also used PubChem for ligands and RCSB-PDB for proteins, consistent with the approach used in the present study, except that the previous study used the method for antiviral drug discovery rather than for microplastic analysis.[16] A study where molecular docking of cellulose obtained from Bacillus species was carried out with cellotetraose followed a similar technique of docking a natural polymer to determine its binding characteristics, similar to the present study in which cellulose was docked with the PET surrogate terephthalic acid. In that case, however, the study was more concerned with enzymatic substrate binding for biofuel production than the capture of microplastics.[17] A study on natural polymers including cellulose and starch as biodegradable alternatives to plastics considered how these materials interacted with plastic components through their structural and molecular properties, which aligns with the present study's finding that cellulose showed the highest binding affinity to PET among all plant polymers tested.[18]

However, one limitation of this research is that instead of employing entire polymer chains, only the surrogates of the microplastics were used in the study, which might not properly represent the surface chemistry of microplastics present in water. Secondly, the docking was done computationally and not through any practical experiments in a wet laboratory environment that could have confirmed the computational findings of this research in the real world. Some recommendations for future research on this topic include conducting some experimental filtration tests using fabric filters impregnated with Moringa seed extract and cellulose or starch derivatives, followed by validation of the computational results through light microscopy and turbidity measurements. In conclusion, this study has provided a promising computational foundation in developing a biodegradable plant-based fabric filter for microplastic removal from water, and future research needs to be done on the use of MOCP and cellulose materials experimentally.

Works Cited

1.Osman, A.I., et al., Microplastic sources, formation, toxicity and remediation: a review. Environ Chem Lett, 2023: p. 1-41.

2.Patil, S., et al., Estimated exposure to microplastics through national and local brands of bottled water in Central India. Environ Monit Assess, 2024. **196**(3): p. 240.

3.Avazpour, S. and M. Noshadi, Enhancing the coagulation process for the removal of microplastics from water by anionic polyacrylamide and natural-based Moringaoleifera. Chemosphere, 2024. **358**: p. 142215.

4.Reza, T., et al., Microplastic Removal in Wastewater Treatment Plants (WWTPs) by Natural Coagulation: A Literature Review. Toxics, 2023. **12**(1).

5.Girish, N., N. Parashar, and S. Hait, Coagulative removal of microplastics from aqueous matrices: Recent progresses and future perspectives. Sci Total Environ, 2023. **899**: p. 165723.

6.Dandesa, B., D.A. Akuma, and E. Alemayehu, Water purification improvement using moringa oleifera seed extract pastes for coagulation follow scoria filtration. Heliyon, 2023. **9**(7): p. e17420.

7.Tang, W., et al., The removal of microplastics from water by coagulation: A comprehensive review. Sci Total Environ, 2022. **851**(Pt 1): p. 158224.

8.Gao, W., et al., Removal of microplastics from water by coagulation of cationic-modified starch: An environmentally friendly solution. Sci Total Environ, 2023. **904**: p. 166787.

9.Lapointe, M., et al., Understanding and Improving Microplastic Removal during Water Treatment: Impact of Coagulation and Flocculation. Environ Sci Technol, 2020. **54**(14): p. 8719-8727.

10.Peydayesh, M., et al., Sustainable Removal of Microplastics and Natural Organic Matter from Water by Coagulation-Flocculation with Protein Amyloid Fibrils. Environ Sci Technol, 2021. **55**(13): p. 8848-8858.

11.Subair, A., et al., Removal of polystyrene microplastics using biochar-based continuous flow fixed-bed column. Environ Sci Pollut Res Int, 2024. **31**(9): p. 13753-13765.

12.Raj, S., B. Mahanty, and S. Hait, Coagulative removal of polystyrene microplastics from aqueous matrices using FeCl(3)-chitosan system: Experimental and artificial neural network modeling. J Hazard Mater, 2024. **468**: p. 133818.

13.Batista, G.S., et al., Removal of Microplastics from Drinking Water by Moringa oleifera Seed: Comparative Performance with Alum in Direct and in-Line Filtration Systems. ACS Omega, 2026. **11**(4): p. 6602-6612.

14.Park, J.W., et al., Removal of microplastics via tannic acid-mediated coagulation and in vitro impact assessment. RSC Adv, 2021. **11**(6): p. 3556-3566.

15.Rath, S., M. Jagadeb, and R. Bhuyan, Molecular docking of bioactive compounds derived from Moringa oleifera with p53 protein in the apoptosis pathway of oral squamous cell carcinoma. Genomics Inform, 2021. **19**(4): p. e46.

16. Souza, H.C.A., et al., Molecular Docking and ADME-TOX Profiling of Moringa oleifera Constituents against SARS-CoV-2. Adv Respir Med, 2023. **91**(6): p. 464-485.
17. Golchha, N.C., et al., Molecular docking analysis of cellulose from Bacillus species with cellotetraose. Bioinformation, 2025. **21**(5): p. 947-951.
18. Ponnusamy, P.G. and S. Mani, Material and Environmental Properties of Natural Polymers and Their Composites for Packaging Applications-A Review. Polymers (Basel), 2022. **14**(19).

Comparative Evaluation of the Physiochemical Properties of Gluten and Gluten-Free Blondies By Isabella Jacobson

Abstract

The gluten-free diet (GFD) is rapidly growing as people desire healthier lifestyles or are gluten-intolerant. Despite growing market trends, gluten-free (GF) products tend to be more expensive and lower in protein and calcium. There is limited research comparing potential GF flours in formulating a GF product that is most similar to a gluten-containing (GC) product. This study examined the similarities in water activity, color, weight, volume, and protein and calcium content across four GF flours (almond, chickpea, coconut, and peanut) to those of all-purpose flour (APF). Almond flour blondies (AFB) were most similar in weight, chickpea flour blondies (CHB) in color and protein content, coconut flour blondies (COB) in volume, and peanut flour blondies (PFB) in water activity and calcium content. There was no clear conclusion as to which flour could create the most similar blondie to the AP blondie (APB), but for commercial purposes, chickpea flour (CHF) is the best choice due to appearance.

Key Words Blondies, Gluten-Free, Almond Flour, Chickpea Flour, Coconut Flour, Peanut Flour

Introduction

Gluten is responsible for the viscoelastic properties of doughs/batters, and textural components of foods. Especially in processed goods like breads and cookies, gluten enhances flavor, gas and moisture retention, and cohesion, among other physiochemical properties. Gluten can cause autoimmune disorders like celiac disease (CD) and dermatitis herpetiformis (DH), which inflame the small intestine and irritate skin, respectively. The only treatment for these diseases is the gluten-free diet (GFD). While gluten is a functional component in baked goods' formation, it is the main cause of autoimmune disorders like CD and DH.

As more people strive to live healthier lifestyles, they may choose to adopt a diet. Manufacturers need to adapt their food formulations to new dietary trends. Specifically, the GFD is growing in influence, as evidenced by the global market for GF products being $8.12B in 2025 and a projected $22.4B in 2035 having CD, roughly 25% of the Western hemisphere's population, medical complications, or believing GFDs have even proven to cause weight gain. To reiterate, the GFD has become popular because of assumptions regarding its health benefits, demonstrating the significance of this study and the scope of potentially impacted people.

There are various availability, cost, and nutritional discrepancies between GC products and their GF counterparts. A study that compared the price, accessibility, and nutrients of GF products found in grocery stores with equivalent GC products discovered that GF commodities are 42% less available in the market than GC items.5 This explains why GF products are more expensive than their GC counterparts, often over +250% the original price. They also found a 69% difference in protein content in GF goods. Further, the GFD lacks calcium, putting people

with CD at risk of osteoporosis, evidenced by 75% having low bone density. Additionally, products containing GF flour exhibit poor cohesive, elastic, and physiochemical properties (e.g., brittle crumb, decreased shelf life, and blandness). This establishes the greater demand for GF foods, especially those most like GC products. In summary, compared to GC foods, GF products demonstrate a trend of being higher in price despite nutritional deficits of calcium and protein and poor physiochemical properties.

It is important to define GF and "blondie" to better comprehend this study. To qualify as GF, food must contain ≤ 20 ppm gluten. Erika Kwee, winner of the 2021 Southern Living Cook of the Year, describes blondies as being like a brownie sans cocoa with crackly tops. Patty Catalano, Food Editor from the Kitchn, with an MS in Food, Nutrition and Culinary Science, believes blondies should be golden, have firm edges, and chewy centers. Overall, blondies should be fudgy with firm crusts, crackled on top, and look golden.

The objective of this research is to measure the qualitative and quantitative attributes of GF blondies (GFB) utilizing four GF flours to identify which is most like the GC blondie (GCB). The research question analyzed is: "Which gluten-free flour can be used to create a Blondie that is most similar physiologically and nutritionally to a gluten-containing blondie?" The objectives are: 1) identify a feasible recipe an average baker can follow to create GFB at home; 2) determine which GF flour produces a blondie most similar in physiological characteristics to the GCB; and 3) determine if the flours with higher nutritional content correspond with the most similar blondies.

Literature Review

Gluten and Celiac Disease

It is beneficial to review existing literature to understand why gluten is significant. Associate Professor of Human Nutrition at the University of Melbourne, Jessica Biesiekierski, with a Ph.D. in non-celiac gluten sensitivity and irritable bowel syndrome, gluten as "the main storage protein of wheat grains. [It] is a complex mixture of hundreds of related but distinct proteins, mainly gliadin and glutenin." Common gluten-containing products include pastas, breads, and beer. In a study Biesiekierski conducted, she explained how gluten is commonly incorporated into foods because it acts as a binding and extending agent to improve overall acceptability by enhancing the products' taste and textural properties. Consequently, a lack of gluten creates challenges in preserving a food's typical characteristics, which is important for incentivizing people to consume the product.

The GFD

Some follow the GFD due to gluten sensitivities including CD and DH, irritating the intestine and skin, respectively. The GFD is commonly associated with living healthier, despite limited evidence, and has increased prominence. Thus, there is an increasing demand for GF products, explaining their higher cost compared to GC foods. In the U.K., Canada, and Austria,

GF products were 159%, 242% and 267% more expensive, respectively. This proves the need for cheaper GF products and raises the question, "*What is incentivizing people to buy them?*" One explanation is the Health Halo Effect, described as "claims concerning just one nutrient, as consumers often make generalizations about the overall health of a product based on one piece of information found on labels." However, these products often lack calcium, iron, magnesium, and protein. Calcium deficiencies contribute to bone fragility/diseases in people diagnosed with CD. The growing popularity of a GFD diet reinforces the need to expand GF options. Brands like Oreo and Kraft, for example, offer GF alternatives and the projected global market value of GF products generally is $13.67B. This demonstrates manufacturers adapting to diet fads to attract more consumers. The growing demand for GF foods despite their higher prices explains why addressing GF options is relevant.

GF Flours

Previous research on GF products analyzed how GF alternatives affected food characteristics including acceptability, water absorption capacity, and color. Peanuts and almonds are high in protein, which has been proven to cause an increase in water absorption since proteins efficiently bind to water molecules, so it was expected that PF and ALF would contain the greatest protein content. Nuts increase the fat content, which decreases product volume because it disrupts the gluten network. PF foods have low gas retention and moisture diffusion ability, contributing to shorter heights. Previous research involving PF also observed increases in density and decreases in moisture content, attributed to PF's lower starch content. There were also observed increases in intensity of brownness, explained by the Maillard Browning Effect. Because the flour color impacts the product's color, and peanuts are darker in color compared to the other nuts in this study, I predict PB will be the darkest.

ALF utilized in a comparison study of CHF and ALF in cupcakes found that there was a decrease in volume of ALF cupcakes, despite almonds' high fat content. Other research involving experimentation with ALF in GF bread found that volume, moisture content, and cohesion was positively correlated. There is a clear discrepancy in ALF's effect on volume. Although this study will not address the effect of various amounts of ALF in a blondie, it will be compared with PF to determine similarity. I hypothesize that AFB and PFB will have greater volumes than other GFB.

CHF generates its own effects on foods, as follows. Based on a separate study, CHF's greater emulsion stability accounts for higher cohesion. However, another study contradicts these results, finding an inverse relationship between protein addition and cohesion. I expected CHB hardness would be greater than other GFB due to low water absorption capacity and moisture content. As mentioned, the original flour color impacts the food's color, explaining why CHF desserts tend to be lighter. CHF and PF are opposites: flour addition results in a direct and inverse relationship with volume, respectively. I hypothesized CHB to be lighter and have a greater volume than PFB.

This study utilized COF, so it is important to review literature on its effect on food. COF contains high amounts of fiber, which has hydroxyl groups that easily bind to water molecules, indicating a greater water absorption capacity. A study conducted in India agrees through its examination of physiochemical properties of muffins with increased COF. The results demonstrated that volume, water absorption capacity, moisture content, and acceptability increased; density decreased. However, separate research discovered that not only cohesiveness increased, but also hardness by adding COF. This is unusual because the higher moisture content associated with COF should cause tenderness.

Gap

There is ample research comparing how GF flour impacts treats such as cupcakes, muffins, cakes, and cookies. However, there is a gap in studies regarding GF bars. I chose to focus on blondies because they are more niche and not abundant in markets compared to other popular bars like brownies. In existing research, GF flours are tested and combined to discover the most successful substitution ratio in a food; the effect of amounts of one GF flour is analyzed; and acceptability between GF flours. This study addresses a gap because it compares several GF flours and to each other and a GC dessert. Wheat flour is often a control group when evaluating the appropriate ratio of flours for an improved product, whereas this study uses APF for control. Studies focused on the acceptability of GF flours in baking without comparing resemblance to its GC counterpart. This is vital as similarity to the original product is a main purchasing/consumption factor. (El Khoury et al., 2018) validates this, acknowledging disparities in GF vs. GC products and a priority of creating new GF products to achieve GC-like characteristics: “it would be more useful to compare the obtained gluten-free ‘dough’ and bread characteristics with those of an actual similar gluten-containing system.” This study addresses this by determining which GF flour will create a blondie most similar in physiological attributes to the one containing gluten. The research seeks to answer the question: Which gluten-free flour can be used to create a Blondie that is most similar physically and nutritionally to a gluten-containing blondie?

Methodology

Blondie Formulation

The overall methodology was based on (Dogruer et al., 2023) which evaluated the effectiveness of chickpea, hazelnut, and carob flours in GF cookie formulations. They used the same recipe for all flours, replacing only the flour-to-flour ratio; a colorimeter measured color; and weight of the dough vs. baked cookie was measured daily to evaluate moisture and baking weight loss.

The recipe used is authored by Nicole Hunn, founder of Gluten Free on a Shoestring, with sixteen years of GF baking experience. Sarah Fennel, a.k.a. Broma Bakery, also authored a GFB recipe. Fennel has over three million followers across social media, authoring over 1,000

recipes, nineteen of which are GF, and one cookbook, *Sweet Tooth*. Conversely, Hunn has published five GF cookbooks; been featured across acclaimed media platforms, including The Kitchn, a top food website in the U.S., and the Today Show. For these reasons, I followed Hunn's recipe, as described below.

The dry ingredients were whisked together. Granulated sugar and chopped unsalted butter were whisked in a saucepan over heat 4 for 2.5 minutes until simmering. Light brown sugar was added and melted. The eggs and vanilla extract were beaten with a hand mixer (Hamilton Beach Brands, Glenn Allen, VA) for 1.5 minutes at speed 2. While beating, the sugar mixture was combined with the eggs. A wooden spoon stirred in the flour blend and chocolate. The batter was poured into a prepared (B&G Foods, Inc, Parsippany, NJ) 18.73cmx18.73cmx3.33cm pan (Target Corporation, Minneapolis, MN). Blondies were placed in a preheated (162.78°C) conventional oven for 50 minutes. ALB were left in the oven for longer after coming out raw, and after 70 minutes they were still underbaked. They cooled on a rack for one hour before being cut into nine pieces and placed in labeled zip-sealed plastic bags according to flour-type (figures 1-2). See figures 3-8 for baking and product visuals. See table 1 for specific amounts and manufacturing information.

To create a control group, APF was substituted in a 1:1 ratio with ALF, CHF, COF, and PF. See Table 2 for further flour information. Nicole Hunn recommends xanthan gum (XG) should be added if the flour doesn't already contain it and should not be used with APF. XG is a hydrocolloid known to strengthen viscosity and enhance gas retention, contributing to higher volume. PB were the only ones to omit XG.

Physical Properties

Three blondies were reserved for quantitative analysis, and six were left for weight measurements. Data was collected within a week of baking. Rubber gloves were worn to minimize contamination.

Color

A colorimeter (Konica Minolta Sensing Americas, Inc., Tokyo, Japan) measured L values39 (figures 5-7) and personal observations were also recorded (Table 3). Color is significant because it influences purchasing decisions. This method is justified because most research evaluating color in foods utilizes a colorimeter, and it is highly accurate. A trial was performed on 2/21/2026 to ensure recipe feasibility; L values were the only data recorded. Most measurements were collected on 2/28/2026.

Weight

Weight was measured alternate days for five days using a food scale (Amazon.com, Inc., Seattle, WA) to determine free water loss. This data can help determine the rate of staling. A new paper towel was placed on the scale after each type of blondie to reduce cross-contamination. The scale was zeroed out to avoid measuring additional weight. The blondies were placed on the

scale, the weights recorded, then placed back in the bag.

Volume

Volume measurements were calculated using Formula 1. To ensure equal length and width, the pan's dimensions were divided by three so they could be cut in nine. The new length and width equaled 6.2cm. The heights were collected using a ruler after one hour because blondie heights decrease upon cooling.

Rheological Properties

Water Activity

A water activity meter (Addium Inc., Pullman, WA) measured water activity44. It was analyzed because it is a common measurement in various other studies with purposes similar to this one. Another researcher asserts that "the most important quality parameter that determines the quality, texture changes, and microbiological shelf life of baked goods is water activity." Triplicates were measured by cutting out a cube from each and placing them in individual containers (figures 8-10). The approach to using a water activity meter and triplicates is based upon the method suggested by (Singh et al., 2026).

Nutritional

The blondies' nutritional composition was formulated using Recipal, a nutrition label software used by successful companies, including Lesser Evil, known for its organic popcorn. The ingredients, including specific brands and amounts used, were inputted (figure 11).

Results

Physical Characteristics

Color was measured qualitatively through descriptions and visuals, and quantitatively through colorimeter measurements. PFB had the darkest color and a strong peanut scent; CHB appeared the most characteristically blondie golden-yellow; APB displayed swirls of chiffon and linen colors; COB looked the driest and had the strongest aroma; and ALB appeared raw and gooey.

The data in Figures 12-14 demonstrates how the blondies' colors differed. On February 21st, the CHB was the most similar in lightness (L) value to the APB, with a difference (formula 2) of 1.14 in average value. On February 28th, the COB had the least difference (1.77) in L values compared to the APB. Because this difference was greater, the CHB were decidedly most similar in color to APB. The ALB contrasted most in L values with a 15.91 and 15.57 difference in average values each day, respectively.

Volume was measured using a ruler to determine the blondies' length, width, and height. APB had an average volume of 138.38cm3. AFB had the greatest average volume (172.98cm3)

which was the most distinct from APB. COB had an average volume of 134.54cm3. They exhibited the closest volume to APB (variation of 3.84). PB had the smallest volume (123.01cm3). To summarize, ALB had the most different volumes and COB had the most similar volumes to APB.

The average weights were collected from 2/27/2026 - 3/3/2026. Across the five days, COB proved to be the outlier and deviated the greatest from APB average weight. These blondies exhibited the smallest weights across the data collection period (figure 15). The variation values were 21.6, 16.2, and 7.6, on each respective day. COB increased in weight as well, suggesting adsorption.47 It differs from the other blondies that are desorbed, evidenced by the weight decrease. The final weight of the APB was 102.00g, a decline from its original weight of 111.8g. CHB had the most similar weight to APB on 2/27/2026 (variation of 0.8). On 3/1/2026,

ALB demonstrated the closest weight to APB (difference of 1.40). On 3/3/2026, PB had the most similar weight (difference of 2.00). Since CHB's difference is the lowest, it is determined that CHB would be considered the closest weight to APB.

Physiochemical Properties

Water activity is an important measurement to compare the shelf-life stability of each blondie. Water activity values above 0.7 indicated there is a good chance of microbial growth and faster spoilage in food.48 APB had an average of 0.760, which was the greatest value compared to the GF blondies. CHB displayed the lowest water activity (0.718) and a difference of 0.042. PB had the lowest variation value (0.04) with a water activity value of 0.756. The ALB and COB water activity measurements were close to the APB, CHB, and PB values. Thus, water activity may not be a significant data point to consider when determining which flour to use to create the most similar GFB to a GCB. Further, because the water activity levels are all above 0.7, these products are subject to microbial growth and staling. See Table 4 for the complete list of recorded water activity values.

Nutrients

To manage the scope of this research, and because of the previously addressed lack of calcium and protein common in GF products, this study only focused on calcium and protein contents. It is also important to note that blondies are still considered desserts and should not be considered "healthy" alternatives to more nutritious whole foods. This research only serves to determine which GFB could provide the most calcium and protein contents. APB have the lowest protein content (6g) and out of the GFB, CHB have the most similar protein content (8g). Protein is highest in the PB (13g), as expected, which is over double the amount in the APB. Meanwhile, ALB and COB have the same amount (9g). Their calcium content does not reflect their similarity, however, as ALF's resulted in the highest calcium (130mg) vs COF's (80g).

Although PB varied the most from the APB in protein content, it contained the same value (70mg) of calcium. CHB had the second highest (100g of calcium). See Table 5 for a broader view of multiple nutrients in each blondie.

Discussion

Peanut Free Blondies (PFB)

Most of the data from this study regarding the PFB matched the results of existing research studies. A comparison study examined the different effects of PF vs. ALF on GF cookies.49 It validated the data pertaining to a greater weight of the ALB because they found that ALF cookies had greater moisture content, which correlates with heaviness. They also agreed that protein content, "L" values, and water activity should be higher in PF products than ALF. PFB had the highest protein content which aligns with (Mkoko et. al., 2025) which evaluated the effect of different percentages of PF in a wheat flour blend. They similarly observed an increase in protein content with the addition of PF, which can be attributed to peanuts being an exceptional source of protein. They also recorded a decrease in weight with more PF, which correlates with this study's findings that PFB had the second lowest weights across the five-day period. A difference was that in Mkoko and her team's research, they observed a direct relationship between the addition of PF and volume, whereas PFBs had a lower volume than the APB in this study. However, multiple pieces of existing research shared these observations of an inverse relationship between PF addition and volume. To elaborate, Shongwe et. al (2022) observed a decrease in bread volume from 1297.0 cm3 to 962.25 cm3 with a greater ratio of PF. Because Mkoko's team used wheat flour, it can be difficult to make a direct comparison. Moreover, PFBs appeared denser as a result of weak gas retention ability, which is common for GF products since they lack a gluten network to provide elasticity and airiness. A different study coincides with this study's observations by noting increases in density with PF products. Another similarity across various studies pertains to color measurements. Due to the Maillard reaction, melanoidin color compounds are released and explain why PB are a darker shade of brown/cream. This reaction occurred in other studies as L values decreased with greater peanut content. Overall, my research findings align with the results of previous literature.

ALF and COB

Analyzing the results of the ALF and COB blondies using data from existing studies, especially a specific comparison study that examined both flours in cupcake formulation, can provide insight into how my findings align or contradict other research. ALF had the greatest volume, which is validated by Gillespie and Ahlborn (2021) and Hopkin et. al's (2022) results. The latter study also aligns with the weight measurements because they recorded an increase in moistness in ALB which has a direct correlation with weight. Since ALB in this study had the greatest weights overall over the three-day period, it is further confirmed that AFB will be the

heaviest. ALF had a relatively high protein content compared with oat bran fiber in other research, which explains the weight increase since protein has a high-water binding capacity.

This study's findings also disagree with results found in Stoin (2018)'s research which involved incorporating 20%, 40% and 60% ALF into GF rice flour blend to make muffins. Stoin noted a decrease in width and length in the ALB, which would indicate a decrease in volume. However, the volume measurements of this research demonstrate that ALB resulted in the largest volume. This could be attributed to the height of the ALB being noticeably greater than the other measured heights. Since Stoin's data did not collect data on height, there is a limitation in the extent to which I can compare these findings to hers because the heights may have similarly increased.

The data collected from the COB disagrees and agrees with existing studies on the impact of COF on GF baked goods. COF had the lowest weight measurements in this study for all three data collection days. This finding does not coincide with Ramya and Anitha's (2020) research which demonstrated an increase in muffin weight with greater COF percentages. COB were the only ones to increase in weight across the testing period, suggesting adsorption, rather than desorption54 in the other blondies. A different study compared COF with APF and discovered that plain cake made with COF had a lower volume than that of gluten containing cake. This matches the data collected in this study. Ramya and Anitha's study (2020) coincides with my results regarding volume because it concluded that there was a direct relationship between coconut content and volume, and COF seems to have likewise increased the blondies' volume. They observed a minimal decrease in protein content and likewise with calcium as it dropped from 23.00mg to 0.31mg. However, this study showed calcium and protein increasing with COF replacement. By comparison, Abayomi and Olalu (2026) proved an increase in protein content of COF cake juxtaposed with APF cake. These results more directly support my data because of the APF control, which Ramya and Anitha did not use. In comparison to APF, COF is a potential GF replacement that will still provide nutritional benefits to those impacted by GFD. In summary, most of the results from this study disagree with the results of existing literature.

CHB

Most of the results in this study aligned with data from existing research on CHB. CHF had the second highest protein content out of all the blondies and Dogurer et. al (2023) showed that CHF products tend to have greater protein content than other GF flour products. CHB having the second highest calcium content is validated by this research, which found significant amounts of calcium in CHF pasta. Dogruer et. al (2023) discovered a decrease in moisture content and weight of CHF pasta, yet my research suggests that CHB have the greatest weights of all GF and APB. In the same study, the recorded "L" value of CHB was approximately 90.60, which is much higher than the calculated "L" values of CHB. However, when comparing the "L" value to other GF flours, they determined that CHF had a lower "L" value, which does align with this study's results. The reason for the lower "L" value is because chickpeas are darker compared to the other nuts used in the GF flours in this study, and smaller "L" values indicate reduced

lightness. A different study agreed, noting that an increase in CHF content led to a darker and browner crust and center color of bread. The qualitative observations regarding the darker appearance of the CHB, PNB, and ALB in comparison to the other GF blondies is supported by research from another study that discovered an increase in protein content seems to correlate with products of lower "L" values and darker hues. This is because the Maillard Reaction is prolonged during baking and increases the browning effect. CHB had a lower volume than APB, which was similarly observed in (Kotsiou et al., 2022) and (Mohammed et al., 2012). This is because the lack of a gluten network prevents efficient trapping of gas in the dough, which expands the blondies to create a lighter, fluffier texture. This explains why GF products tend to be denser than their gluten-containing counterparts.

Retrodegradation

Starch retrogradation explains the observed trend in blondie weights across the five-day testing period. It is a primary factor impacting the hardness, texture, flavor, scent, and shelf-life of products as it causes staling. All the blondies except for the COB demonstrated a decrease in weight. This is indicative of retrodegradation because the blondies experienced desorption as the free water was lost and absorbed into the atmosphere, thus increasing the blondies' hardness. Regarding the COB, since they increased in weight, retrodegradation occurred to a smaller extent since these absorbed. This would make the blondies denser after sitting out, as opposed to the other blondies which became drier and crumblier. When compared to research on retrodegradation in GF noodles, the texture was altered and firmness increased because of starch recrystallization,60 aligning with my findings. They also observed a lower "L" value in the retrograded noodles. ALB exhibited the smallest "L" value and had the greatest weight on average, which may signify high rates of retrodegradation. In contrast, COB had the highest "L" value and the smallest weight, demonstrating lower rates of retrodegradation.

Implications

After completing this study, it is crucial to state its significance to existing research on the utility of and differences between GF flours used to produce baked goods. One objective of this research was to find a suitable recipe that home bakers can follow to create their own GFB. This would allow home bakers to avoid paying higher prices for GF foods, specifically blondies. The recipe from glutenfreeonashoestring.com is a feasible option. Another goal of this study is to identify the GF flour that produced a blondie most physiologically like the control blondie. Based on the data, there is no clear answer regarding water activity, PB displayed the most similar results to the APB; CHB were the closest in "L" value, and thus overall color; volumes of the APB and COB were most alike; and ALB were most similar in weight to the control group. When focusing on nutritional content, PB contained the same calcium content and CHB had the most similar protein content. The results are scattered amongst all four tested GF flours, meaning that which GF flour to use is dependent on what the consumer prefers and believes is the most important factor. Product appearance is a main reason for customers to purchase a

product. Thus, if buyers understand what a GCB is like, they may be drawn to GFB that appear similar. Therefore, CHF may be the best option for manufacturers and product developers to use.

Another objective of this study was to determine whether GF flours with the greatest protein and calcium contents would also be most similar to GCB. There is a specific focus on protein and calcium due to the general nutrient deficiency of these in the GFD, as previously established in this paper. The highest calcium content and protein content were in the ALB and PB, respectively. However, the ALB contained the second highest protein content, and PB had the lowest calcium content. Thus, the ALB would be the "most nutritional", considering that these desserts are not healthier substitutes for other food. Therefore, Nicole Hunn's recipe would be a suitable recipe for home bakers to follow who are looking to make their own GF products to avoid some of the higher prices of GF foods in stores, food product developers looking to maximize profits from blondies should use CHF, and although ALB proved have the most beneficial nutritional content for someone on a GFD, it wasn't the most similar to the control group.

Limitations

Although this study reached important conclusions about GFs used to produce GFB, there are limitations that challenge the accuracy of the research. First, the GF flours were used in the same recipe in a 1:1 ratio despite the disagreements over interchangeability of the flours. Food Network suggests using weight, not volume to substitute between GF and GC flours. Rachel Ciordas, author of GF recipes for over a decade, agrees with this assertion, claiming that using "cup-for-cup" conversions is not the most appropriate strategy. She adds that nut flours can commonly be substituted for each other, but because they add grease to the final product, they should not be used interchangeably with other flour types. Similarly, PF can be used in a 1:1 ratio with APF, but it is suggested that the amounts of other ingredients in the recipe be adjusted to minimize greasiness. Further, another article states that COF should not be used in a 1:1 ratio with any GC flour. Kasey Trenum, recipe creator and blogger, asserts that ALF and APF can be substituted 1:1, but recommends bakers use 1/4 or 1/3 cup COF for every 1 cup of ALF. This indicates that COF cannot be used in equal amounts with APF. However, CHF can be used in a 1:1 ratio with ALF, and thus APF as well.

Furthermore, there were very limited studies performed measuring the water activity and calcium content of the GF flours in this recipe. This demonstrates a gap in the research, yet it also makes discussing the results for validation and understanding and how they align with existing literature challenging. Because of this, there is minimal discussion regarding calcium content and water activity. Future research could examine the various water activities across GF foods to evaluate shelf-life stability and identify longer-lasting products. Another extension of this study could involve a sensory evaluation of GFB to determine consumer acceptance of products. Such a study, for example, could evaluate whether the blondies that were most like the

APB were the highest scored by sensory panelists. This research may also be replicated with other GF flours, such as oat and sorghum.

Conclusion

This research identified a feasible recipe for people with gluten-intolerance to follow to bake blondies to avoid the high prices of GF baked goods. Regarding the research's second objective, the results did not identify a definitive answer as to which GF flour could be used to produce the most similar GFB to a GCB. COB was most similar in volume; CHB in "L" value, weight, and protein content; and PF in water activity and calcium content. The data also suggests that the blondies with the highest protein and calcium content did not correlate to the most similar blondie in appearance to GCB. Finally, this study has identified a potential commercial use of CHB because the flour produces a similar appearance to blondies containing gluten that would be more likely to encourage customers to purchase the product. This would increase the prevalence of GFB in grocery stores and provide a broader variety of options for gluten-intolerant individuals.

Works Cited

Samuel Solanke, Abayomi, and A. F. Olalu. “Consumer Acceptability of The Sensory and Proximate Properties of a Plain Cake Produce from Flour Blends (Oat, Tigernut and Coconut).” International Journal of Innovative Research in Technology, vol. 12, no. 10, 2026, pp. 5006–5012. www.researchgate.net/publication/402868924_Consumer_Acceptability_of_The_Sensory. Accessed 17 October 2025.

Absorbing Capacity - An Overview | ScienceDirect Topics. (n.d.). https://www.sciencedirect.com/topics/engineering/absorbing-capacity. Accessed November 1, 2025.

Adeboye, A. S., et al. “Towards the Development of Peanut-Wheat Flour Composite Dough: Influence of Reduced-Fat Peanut Flour on Bread Quality.” Journal of Food Processing and Preservation, vol. 42, no. 1, Wiley, Jan. 2018, p. e13385, https://doi.org/10.1111/jfpp.13385. Accessed 17 October 2025.

Aditya, D. S., et al. “Polymer-Clay Nanocomposites for Dyes Abatement.” 2024. www.sciencedirect.com/science/chapter/referencework/abs/pii/B9780323954860000156. Accessed 16 October 2025.

Görgülü, Ahmet. “Real-Time Quality Analysis of Baked Goods Using Advanced Technologies.” Journal of Food Engineering, vol. 388, no. 112359, Elsevier BV, Mar. 2024, p. Not Available, https://doi.org/10.1016/j.jfoodeng.2024.112359. Accessed October 25, 2025.

An, Ni, et al. “Adsorption Characteristics of Rocks and Soils, and Their Potential for Mitigating the Environmental Impact of Underground Coal Gasification Technology: A Review.” Journal of Environmental Management, vol. 305, no. 114390, Elsevier BV, Mar. 2022, p. 114390, https://doi.org/10.1016/j.jenvman.2021.114390. Accessed November 3, 2025.

Arya, Shalini S., et al. “Peanuts as Functional Food: A Review.” Journal of Food Science and Technology, vol. 53, no. 1, Springer Science and Business Media LLC, Jan. 2016, pp. 31–41, https://doi.org/10.1007/s13197-015-2007-9. Accessed November 13, 2025.

Belorio, Mayara, et al., “Influence of Flour Particle Size Distribution on the Quality of Maize Gluten-Free Cookies.” Foods (Basel, Switzerland), vol. 8, no. 2, MDPI AG, Feb. 2019, p. E83, https://doi.org/10.3390/foods8020083. Accessed November 20, 2025.

Biesiekierski, Jessica R. “What Is Gluten?” Journal of Gastroenterology and Hepatology, vol. 32, no. S1, Wiley, Mar. 2017, pp. 78–81, https://doi.org/10.1111/jgh.13703. Accessed October 13, 2025.

Cardo, Aner, et al. “Nutritional Imbalances in Adult Celiac Patients Following a Gluten-Free Diet.” Nutrients, vol. 13, no. 8, MDPI AG, Aug. 2021, p. 2877, https://doi.org/10.3390/nu13082877. Accessed November 24, 2025.

Celiac.org/gluten-free-living/what-is-gluten/sources-of-gluten/. Accessed October 5, 2025.

Christoph, Mary J., et al. “Who Values Gluten-Free? Dietary Intake, Behaviors, and Sociodemographic Characteristics of Young Adults Who Value Gluten-Free Food.” Journal of the Academy of Nutrition and Dietetics, vol. 118, no. 8, Elsevier BV,

Aug. 2018, pp. 1389–1398, https://doi.org/10.1016/j.jand.2018.04.007. Accessed September 24, 2025.

Ciordas, Rachel. “Gluten Free Flour Guide.” Next Level Gluten Free, Aug. 2022. rachelciordas.substack.com/p/gluten-free-flour-guide. Accessed September 15, 2025.

Collin, Pekka, et al. “Dermatitis Herpetiformis: A Cutaneous Manifestation of Coeliac Disease.” Annals of Medicine, vol. 49, no. 1, Informa UK Limited, Jan. 2017, pp. 23–31, https://doi.org/10.1080/07853890.2016.1222450. Accessed November 3, 2025.

Dat, Lai Quoc. “Functional Properties and Influences of Coconut Flour on Texture of Dough and Cookies.” Vietnam Journal of Science and Technology, vol. 55, no. 5A, Publishing House for Science and Technology, Vietnam Academy of Science and Technology (Publications), Mar. 2018, p. 100, https://doi.org/10.15625/2525-2518/55/5a/12184. Accessed December 1, 2025.

Definition and Clinical Manifestations - Celiac Disease Center at Columbia University Medical Center. 2018. https://celiacdiseasecenter.columbia.edu/celiac-disease/definition-and-clinical-manifestations. Accessed October 26, 2025.

Celiac Disease Foundation. (2018). Sources of Gluten | Celiac Disease Foundation. Celiac Disease Foundation; Celiac. https://celiac.org/gluten-free-living/what-is-gluten/sources-of-gluten. Accessed September 25, 2025.

“Chickpea Seed Flours Improve the Nutritional and the Antioxidant Profiles of Traditional Shortbread Biscuits: Effects of In Vitro Gastrointestinal Digestion.” Antioxidants, vol. 13, no. 1, pp. 118–118.

Delgado-Andrade, Cristina, et al, “Chickpea Seed Flours Improve the Nutritional and the Antioxidant Profiles of Traditional Shortbread Biscuits: Effects of in Vitro Gastrointestinal Digestion.” Antioxidants (Basel, Switzerland), vol. 13, no. 1, MDPI AG, Jan. 2024, p. 118, https://doi.org/10.3390/antiox13010118. Accessed December 13, 2025.

Demissie, Yodit Fikadu, et al. “Effect of Blending Ratio and Baking Temperature on Nutritional Quality and Sensory Acceptability of Bread Developed from the Blends of Wheat, Sweet Lupine Flour and Moringa Leaf Powder.” Food Chemistry Advances, vol. 11, no. 101305, Elsevier BV, June 2026, p. 101305, https://doi.org/10.1016/j.focha.2026.101305. Accessed June 15, 2026.

Dennis, Melinda, et al., “Nutritional Considerations of the Gluten-Free Diet”. Gastroenterology Clinics of North America, vol. 48, no. 1, Elsevier BV, Mar. 2019, pp. 53–72, https://doi.org/10.1016/j.gtc.2018.09.002. Accessed October 17, 2025.

Dermatitis Herpetiformis (For Health Care Professionals) | NIDDK. (2014, February). National Institute of Diabetes and Digestive and Kidney Diseases. https://www.niddk.nih.gov/health-information/professionals/clinical-tools-patient-management/digestive-diseases/dermatitisherpetiformis. Accessed December 12, 2025.

Doggwiler, M. “The Maillard Reaction: The Science behind Flavour and Colour in Foods and Beverages.” Ragus.
2024, https://www.ragus.co.uk/the-maillard-reaction-the-science-behind-flavour-and-colour-in-foods-and-beverages/. Accessed January 3, 2026.

Dogruer, I., et al. “Formulation of Gluten-Free Cookies Utilizing Chickpea, Carob, And Hazelnut Flours Through Mixture Design.” Foods, vol. 12, no. 19, 2023, p. 3689, https://doi.org/10.3390/foods12193689. Accessed January 13, 2026.

Dubat, Arnaud. “Water Absorption Capacity of Flour.” KPM Analytics, 27 May 2026, www.kpmanalytics.com/blog/water-absorption-capacity-of-flour. Accessed January 14, 2026.

El Khoury, D., et al. “A Review on the Gluten-Free Diet: Technological and Nutritional Challenges.” Nutrients, vol. 10, no. 10, 2018, p. 1410, https://doi.org/10.3390/nu10101410. Accessed December 15, 2025.

Emily, Clive, et al. “Understanding the Changes in Brain Activation When Viewing Products with Differences in Attractiveness.” Neurology International, vol. 16, no. 5, 2024, pp. 918–32, https://doi.org/10.3390/neurolint16050069. Accessed February 10, 2026.

Estévez, V., et al. “Persistent Barriers of the Gluten-Free Basic Food Basket: Availability, Cost, And Nutritional Composition Assessment.” Nutrients, vol. 16, no. 6, 2024, p. 885, https://doi.org/10.3390/nu16060885. Accessed January 17, 2026.

Unknown. “Gluten-Free Means What It Says.” Webpage. FDA Consumer Updates. https://www.fda.gov/consumers/consumer-updates/gluten-free-means-what-it-says. Accessed December 31, 2025.

Fetriyuna Fetriyuna, et al. “The Role of Xanthan Gum and Hydroxypropyl Methylcellulose in Gluten‐Free Bread: A Study of Physical Characteristics, Texture, And Nutrition.” Food Science & Nutrition, vol. 14, no. 1, p. 71292–92, https://doi.org/10.1002/fsn3.71292. Accessed January 7, 2026.

Gillespie, Rachel, and Gene J. Ahlborn. “Mechanical, Sensory, And Consumer Evaluation of Ketogenic, Gluten-Free Breads.” Food Science & Nutrition, vol. 9, no. 6, 2021, pp. 3327–35, https://doi.org/10.1002/fsn3.2308. Accessed January 20, 2026.

Unknown, Gluten-free Food Market. “Industry Analysis, By Type (Baby Food, Pastas & Pizzas.” Snacks & RTE Products, Bakery Products, and Condiments & Dressings), By Distribution Channel (Supermarkets/Hypermarkets, Convenience Stores, Specialty Stores, Drugstores & Pharmacies, And Online Retails), and Regional Forecast.” https://www.fortunebusinessinsights.com/industry-reports/gluten-free-food-market-100188. Accessed July 13, 2026.

Gorgitano, M.T., and V. Sodano. “Gluten-Free Products: From Dietary Necessity to Premium Price Extraction Tool.” Nutrients, vol. 11, no. 9, 2019, https://doi.org/10.3390/nu11091997. Accessed December 17, 2025.

Granato, D., and L.de S.N. Ellendersen. “Almond and Peanut Flours Supplemented with Iron as Potential Ingredients to Develop Gluten-Free Cookies.” Ciência E Tecnologia De

Alimentos, vol. 29, no. 2, 2009, pp. 395–400, https://doi.org/10.1590/s0101-20612009000200026. Accessed December 20, 2025.
Hartley, Georgina. “Guide to Using Chickpea Flour in Gluten-Free Baking.” From the Larder, 26 Sept. 2024, https://fromthelarder.co.uk/chickpea-flour/. Accessed September 22, 2025.
Higa, Erin. “Peanut Flour vs All-Purpose Flour: The Final Verdict.” The Coconut Mama, 6 Apr. 2023, https://thecoconutmama.com/peanut-flour-vs-all-purpose-flour/. Accessed September 21, 2025.
Hinds, M. J. “Physical Properties of Muffins Containing Peanut Flour and Peanut Butter.” Peanut Science, vol. 30, no. 1, Jan. 2003, pp. 67–73, https://doi.org/10.3146/pnut.30.1.0014. Accessed January 3, 2026.
Hopkin, Lauren, et al. “Influence of Almond and Coconut Flours on Ketogenic, Gluten-Free Cupcakes.” Food Chemistry: X, vol. 13, Mar. 2022, p. 100182, https://doi.org/10.1016/j.fochx.2021.100182. Accessed November 28, 2025.
Horstmann, Stefan, et al. “Starch Characteristics Linked to Gluten-Free Products.” Foods, vol. 6, no. 4, Apr. 2017, p. 29, https://doi.org/10.3390/foods6040029. Accessed January 13, 2026.
Hossain, Shakhawat, et al. “Incorporation of Coconut Flour in Plain Cake and Investigation of the Effect of Sugar and Baking Powder on Its Baking Quality.” International Journal of Nutrition and Food Sciences, vol. 5, no. 5 1, Jan. 2016, pp. 31–38, https://www.researchgate.net/publication/316885049_Incorporation_of_Coconut_Flour_in_Plain_Cake_and_Investigation_of_the_Effect_of_Sugar_and_Baking_Powder_on_Its_Baking_Quality. Accessed January 3, 2026.
Hunn, Nicole. “Classic Gluten Free Blondies.” Gluten Free on a Shoestring, 5 Dec. 2024, https://glutenfreeonashoestring.com/gluten-free-blondies/. Accessed 1 May 2026.
Hussein, Ahmed, et al. “Development and Evaluation of Gluten Free Pasta from Brown Rice, Quinoa, and Chickpeas for Celiac Disease Patients Open.” June 2025, https://doi.org/10.1038/s41598-025-06516-6. Accessed October 3, 2025.
Jamieson, Jennifer A., et al. “Folate Content and Chemical Composition of Commercially Available Gluten-Free Flour Alternatives.” Plant Foods for Human Nutrition, vol. 75, no. 3, July 2020, pp. 337–43, https://doi.org/10.1007/s11130-020-00833-z. Accessed November 3, 2025.
Kilcast, David, and Persis Subramaniam. Food and Beverage Stability and Shelf Life. Woodhead Pub, 2011.
Kotsiou, Kali, et al. “Physicochemical and Functional Aspects of Composite Wheat-Roasted Chickpea Flours in Relation to Dough Rheology, Bread Quality and Staling Phenomena.” Food Hydrocolloids, vol. 124, Mar. 2022, p. 107322, https://doi.org/10.1016/j.foodhyd.2021.107322. Accessed November 3, 2025.
Lamacchia, Carmela, et al. “Cereal-Based Gluten-Free Food: How to Reconcile Nutritional and Technological Properties of Wheat Proteins with Safety for Celiac Disease Patients.”

Nutrients, vol. 6, no. 2, Jan. 2014, pp. 575–90, https://doi.org/10.3390/nu6020575. Accessed December 3, 2025.

Lerner, Aaron, et al. “Navigating the Gluten-Free Boom: The Dark Side of Gluten Free Diet.” Frontiers in Pediatrics, vol. 7, Oct. 2019, https://doi.org/10.3389/fped.2019.00414. Accessed October 23, 2025.

Lessenden, Rachel. “Best Chickpea Flour Substitutes (and How to Use Them).” Health My Lifestyle, 27 Mar. 2023, https://healthmylifestyle.com/chickpea-flour-substitutes/. Accessed 24 Apr. 2026.

Li, Ye, et al. “The Role of Gluten in Food Products and Dietary Restriction: Exploring the Potential for Restoring Immune Tolerance.” Foods, vol. 12, no. 22, Multidisciplinary Digital Publishing Institute, Nov. 2023, pp. 4179–79, https://doi.org/10.3390/foods12224179. Accessed September 10, 2025.

Litchford, April, et al. “Increasing Access to Acceptable and Affordable Gluten-Free Baked Goods.” Journal of Human Sciences and Extension, vol. 12, no. 2, Mississippi State University, Sept. 2024, https://doi.org/10.55533/2325-5226.1397. Accessed September 20, 2025.

Liu, Rong, et al. “Retrogradation Behavior of Starch Dough Prepared from Damaged Cassava Starch and Its Application in Functional Gluten-Free Noodles.” International Journal of Biological Macromolecules, vol. 236, Elsevier BV, Mar. 2023, pp. 123996–96, https://doi.org/10.1016/j.ijbiomac.2023.123996. Accessed December 10, 2025.

Lucendo, Alfredo J., and Alvaro García-Manzanares. “Bone Mineral Density in Adult Coeliac Disease: An Updated Review.” Revista Española De Enfermedades Digestivas, vol. 105, no. 3, Mar. 2013, pp. 154–62, https://doi.org/10.4321/s1130-01082013000300006. Accessed January 16, 2026.

Ly, Bao Chau K., et al. “Research Techniques Made Simple: Cutaneous Colorimetry: A Reliable Technique for Objective Skin Color Measurement.” Journal of Investigative Dermatology, vol. 140, no. 1, Jan. 2020, pp. 3–12.e1, https://doi.org/10.1016/j.jid.2019.11.003. Accessed February 1, 2026.

McCann, James. “Food Sick: Busting the Health Halo.” Veri.Co, 2016, https://www.veri.co/learn/health-halo?srsltid=AfmBOor4ZlMGOnZij3hA3z4Vl3Lv-jfARa3sfwRdumT3mulPN4DCY6bT. Accessed 1 Apr. 2026.

McGruther, Jenny. “Coconut Flour: Baking Tips, Substitutions and Recipes.” Nourished Kitchen, 22 Jan. 2019, https://nourishedkitchen.com/baking-with-coconut-flour/. Accessed October 20, 2025.

Missbach, Benjamin, et al. “Gluten-Free Food Database: The Nutritional Quality and Cost of Packaged Gluten-Free Foods.” PeerJ, vol. 3, Oct. 2015, p. e1337, https://doi.org/10.7717/peerj.1337. Accessed November 3, 2025.

Mohammed, I., et al. “Effects of Chickpea Flour on Wheat Pasting Properties and Bread Making Quality.” Journal of Food Science and Technology, vol. 51, no. 9, May 2012, pp. 1902–10, https://doi.org/10.1007/s13197-012-0733-9. Accessed October 13, 2025.

Moldovan, Camelia, et al. “Sensory and Physico-Chemical Characteristics of Muffins Obtained from Non-Conventional Aglutenic Flours.” https://animalsciencejournal.usamv.ro/pdf/2023/issue_1/Art59.pdf. Accessed December 3, 2025.

Moss, Melinda M., et al. “Optimization of Flour-Replacing Ingredients for Low-Carbohydrate, Gluten-Free Muffins via a Mixture Design with Complete Sucrose Substitution by D-Allulose or D-Tagatose.” LWT, vol. 167, Sept. 2022, p. 113779, https://doi.org/10.1016/j.lwt.2022.113779. Accessed October 16, 2025.

Mystkowska, Iwona, et al. “Gluten-Containing Flours and Gluten-Free Flours as a Source of Calcium, Magnesium, Iron and Zinc.” Scientific Reports, vol. 14, no. 1, Nature Portfolio, June 2024, https://doi.org/10.1038/s41598-024-65530-2. Accessed September 15, 2025.

Narges Allahyari, et al. “The Effect of Different Formulations and Baking Conditions on the Physicochemical Properties of Gluten-Free Cakes.” Applied Food Research, Elsevier BV, Aug. 2025, pp. 101301–01, https://doi.org/10.1016/j.afres.2025.101301. Accessed 24 Sept. 2025.

National Institute of Diabetes and Digestive and Kidney Diseases. “Definition & Facts for Celiac Disease.” National Institute of Diabetes and Digestive and Kidney Diseases, 17 Oct. 2019, https://www.niddk.nih.gov/health-information/digestive-diseases/celiac-disease/definition-facts. Accessed January 3, 2025.

Niland, Benjamin, and Brooks D Cash. “Health Benefits and Adverse Effects of a Gluten-Free Diet in Non–Celiac Disease Patients.” Gastroenterology & Hepatology, vol. 14, no. 2, Feb. 2018, p. 82, https://pmc.ncbi.nlm.nih.gov/articles/PMC5866307/#sec11. Accessed November 15, 2025.

Nurul Atiqah Sani, et al. “Effects of Temperature and Airflow on Volume Development During Baking and Its Influence on Quality.” Journal of Engineering Science and Technology, June 2014, vol. 9, No. 3 (2014) 303 – 313. https://jestec.taylors.edu.my/Vol%209%20Issue%203%20June%2014/Volume%20(9)%20Issue%20(3)%20303-313.pdf. Accessed January 5, 2026.

Prieto-Vázquez del Mercado, Pavel, et al. “Protein Ingredients in Bread: Technological, Textural and Health Implications.” Foods, vol. 11, no. 16, Aug. 2022, p. 2399, https://doi.org/10.3390/foods11162399. Accessed January 2, 2026.

Raghavendra, S. N., et al. “Dietary Fiber from Coconut Residue: Effects of Different Treatments and Particle Size on the Hydration Properties.” European Food Research and Technology, vol. 218, no. 6, May 2004, pp. 563–67, https://doi.org/10.1007/s00217-004-0889-2. Accessed November 1, 2025.

Ramya, H. N., and S. Anitha. “Development of Muffins from Wheat Flour and Coconut Flour Using Honey as a Sweetener.” International Journal of Current Microbiology and Applied Sciences, vol. 9, no. 7, July 2020, pp. 2231–40, https://doi.org/10.20546/ijcmas.2020.907.260. Accessed October 25, 2025.

Saha, Dipjyoti, and Suvendu Bhattacharya. “Hydrocolloids as Thickening and Gelling Agents in Food: A Critical Review.” Journal of Food Science and Technology, vol. 47, no. 6, Nov. 2010, pp. 587–97, https://doi.org/10.1007/s13197-010-0162-6. Accessed January 10, 2026.

Sapone, Anna, et al. “Spectrum of Gluten-Related Disorders: Consensus on New Nomenclature and Classification.” BMC Medicine, vol. 10, no. 1, Feb. 2012, https://doi.org/10.1186/1741-7015-10-13. Accessed September 23, 2025.

“Sensory Analysis - an Overview | ScienceDirect Topics.” Sciencedirect.Com, 2009, https://www.sciencedirect.com/topics/medicine-and-dentistry/sensory-analysis. Accessed October 23, 2025.

Shaheen, Sherif, et al. “Dissecting Dietary Melanoidins: Formation Mechanisms, Gut Interactions and Functional Properties.” Critical Reviews in Food Science and Nutrition, vol. 62, no. 32, Taylor & Francis, June 2021, pp. 8954–71, https://doi.org/10.1080/10408398.2021.1937509. Accessed November 3, 2025.

Shaw, Henry. “Water Activity and Its Role in Food Preservation (March 2025) | UC Agriculture and Natural Resources.” UC Agriculture and Natural Resources, 7 Apr. 2025, https://ucanr.edu/program/uc-master-food-preserver-program/article/water-activity-and-its-role-food-preservation. Accessed January 7, 2026.

Shongwe, Sihle G., et al. “Dough Rheology and Physicochemical and Sensory Properties of Wheat–Peanut Composite Flour Bread.” Legume Science, Feb. 2022, https://doi.org/10.1002/leg3.138. Accessed December 28, 2025.

Singh, Arshdeep, et al. “Validation of High Milk Protein Chocolate Chip Cookie Baking Process and Thermal Death Inactivation Parameters of a 4-Serovar Salmonella Cocktail.” Journal of Food Protection, Elsevier, Apr. 2026, p. 100791, https://doi.org/10.1016/j.jfp.2026.100791. Accessed May 26, 2026.

Skendi, Adriana, et al. “High Protein Substitutes for Gluten in Gluten-Free Bread.” Foods, vol. 10, no. 9, Aug. 2021, p. 1997, https://doi.org/10.3390/foods10091997. Accessed December 19, 2025.

Tamanna, Nahid, and Niaz Mahmood. “Food Processing and Maillard Reaction Products: Effect on Human Health and Nutrition.” International Journal of Food Science, vol. 2015, no. 526762, Jan. 2015, pp. 1–6, https://doi.org/10.1155/2015/526762. Accessed November 13, 2025.

Tengetile Mkoko, et al. “Dough Rheology, Physicochemical and Sensory Qualities of Muffins Made from Wheat–Partially Defatted Peanut Composite Flour.” International Journal of Food Science, vol. 2025, no. 1, Hindawi Publishing Corporation, Jan. 2025, pp. 6635106–06, https://doi.org/10.1155/ijfo/6635106. Accessed 25 Nov. 2025.

Testronix, and M. Rohit. “Benefits of Using Colorimeter for Colour Measurement and Quality Control.” Testronix Instruments, 7 Sept. 2023, https://www.testronixinstruments.com/blog/benefits-of-using-colorimeter-for-colour-mea

surement-and-quality-control/?srsltid=AfmBOopC7kbP0hHsFlfq4edJzt5gB52_UpEJERVJOxrmmWQwRM0kZibc. Accessed December 23, 2025.

Torra, Marta, et al. “Chickpea and Chestnut Flours as Non-Gluten Alternatives in Cookies.” Foods, vol. 10, no. 5, Apr. 2021, p. 911, https://doi.org/10.3390/foods10050911. Accessed September 20, 2025.

Toth, Marcell, et al. “Consumers’ Acceptance, Satisfaction in Consuming Gluten-Free Bread: A Market Survey Approach.” International Journal of Celiac Disease, vol. 8, no. 2, June 2020, pp. 44–49, https://doi.org/10.12691/ijcd-8-2-1. Accessed April 1, 2026.

Trenum, Kasey. “Almond Flour Substitutes: Delicious and Healthy Alternatives.” Kasey Trenum, 13 May 2024, https://kaseytrenum.com/almond-flour-substitutes-delicious-and-healthy-alternatives/. Accessed April 24, 2026.

Vici, Giorgia, et al. “Gluten Free Diet and Nutrient Deficiencies: A Review.” Clinical Nutrition, vol. 35, no. 6, Dec. 2016, pp. 1236–41, https://doi.org/10.1016/j.clnu.2016.05.002. Accessed September 15, 2025.

Vinod, B, et al. “Chickpea as a Promising Ingredient Substitute in Gluten-Free Bread Making: An Overview of Technological and Nutritional Benefits.” Food Chemistry Advances, vol. 3, Elsevier BV, Dec. 2023, pp. 100473–73, https://doi.org/10.1016/j.focha.2023.100473. Accessed November 8, 2025.

Wu, Shao-Chi, et al. “The Effect of Tamarind Seed Gum on the Qualities of Gluten-Free Cakes.” Processes, vol. 8, no. 3, Mar. 2020, p. 318, https://doi.org/10.3390/pr8030318. Accessed January 3, 2026.

Yemmireddy, V. K., et al. “Physico-Chemical Properties of Pancakes Made from an Instant Mix Containing Different Levels of Peanut (Arachis Hypogaea) Flour.” Peanut Science, vol. 40, no. 2, July 2013, pp. 142–48, https://doi.org/10.3146/ps13-01.1. Accessed December 3, 2025.

Investigating The Effects of Centroid Motion on Exoplanet Detection Using TESS Light Curve Data By Advik Sharma

Abstract

Transit photometry can be utilised for the detection and classification of exoplanets, with information such as transit depth and centroid motion being vital for differentiating between confirmed exoplanetary detections and false positives such as eclipsing binaries. This highly common false positive is usually classified through the analysis of centroid movement - the estimated center of a star gauged by light intensity - and how it varies during the entire transit. This project investigates the extent to which centroid motion information affects the classification of confirmed exoplanets. This was achieved through an ablation study by training three 1-Dimensional Convolutional Neural Networks (1D CNNs) to prove that centroid movement provides a quantifiable pattern that the model can pick up in the first place, utilising training data sourced from the Mikulski Archive for Space Telescopes (MAST). The results of this study suggest that centroid movement contributed little to the classification of exoplanets, as seen when the model with both flux information and centroid movement information performed nearly identically to the model with only flux intensity information. However, centroid movement does still provide a valid pattern that the model can identify and classify upon, as the model with solely centroid movement data still performed better than simply guessing randomly, albeit rather poorly. This may be attributed to the limitations in terms of the number of examples present for the model to train on, a limitation present as the TESS TOI catalog did not have enough false positives to work with.

1 | Background

Ever since its inception, transit detection has played a pivotal role in the field of exoplanet detection, particularly due to the various exoplanet characteristics that can be reasonably inferred from that single event (Charbonneau et al.; Shallue and Vanderburg). Moreover, the transit method is often combined with alternative methods such as radial velocity, not only allowing for a valid confirmation of a positive reading but also providing more extensive information on the exoplanetary system of choice, such as the minimum mass of the exoplanet (Shallue and Vanderburg). However, when other detection methods are unavailable, it is undoubtedly easy to fall into the category of false positives - falsified detections of exoplanets - the most common of which are eclipsing binaries.

Transit Photometry

This method of exoplanet detection utilises the apparent brightness of a star, more formally known as the observed flux intensity of a star, in an exoplanetary system. When a planet passes in between this star relative to the Earth's viewing perspective, a dip or reduction in the observed flux is recorded up until that said planet has stopped eclipsing the star that lies before it.

The change in observed flux - that is, the difference between the observed flux before and after an eclipse - can be used in order to determine the size of the planet as well as its orbital distance from its star given that the planet is not any smaller than our own planet Earth. This can be done through the following formula (Ar; Shallue and Vanderburg):

$$\frac{\Delta F}{F} = \left(\frac{R_P}{R_S}\right)^2 \tag{3}$$

where R_P represents the planet's radius and R_S represents the star's radius. This is also commonly referred to as the transit depth of an exoplanet. Note that, in general cases, it is R_S that is calculated using this equation, as R_P is generally found through alternative methods of exoplanet detection such as radial velocity imaging.

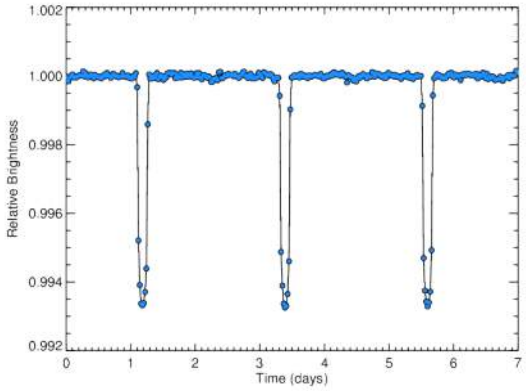


Fig. 1 | The discovery of WASP-14b through Transit Photometry (Levrard et al.). As you can see, the dips in flux are quite evident.

In order to detect these light dips, readings are taken continuously between small intervals known as the 'cadence', which may vary with apparatus and methodology although studies try to minimise the time between each cadence (McDonald and Veras). In other words, the interval between readings is minimised to obtain more precise results pertaining to the start (ingress) and end (egress) times of the observed eclipse. The cadence of the transit photometry method, as well as any captured noise such as stellar or granular oscillation noise (Oshagh), determine the cleanliness of the data captured.

Eclipsing binaries

Whereas transit photometry can yield plenty of information for exoplanet classification, readings resulting in a dip in flux may occur irrespective of an exoplanet eclipsing its parent star. Known as False Positives (FPs), one of the most common examples of this is eclipsing binaries (EBs), in which one star eclipses another when viewed perpendicular to their relative orbits. When one star eclipses another, a decrease in the observed flux, ΔF, is recorded due to either

star’s luminosity being blocked by the one in front of it relative to the viewpoint of Earth (Almenara and Deeg).

Consider an example of two eclipsing stars, as shown in the figure below. The star in the warmer tone is at a temperature of 8,000 Kelvin, whereas the star in the cooler/bluer tone is at a temperature of 6,000 Kelvin. Assume that both stars have the same radius, surface area, mass and orbital (circular) radius.

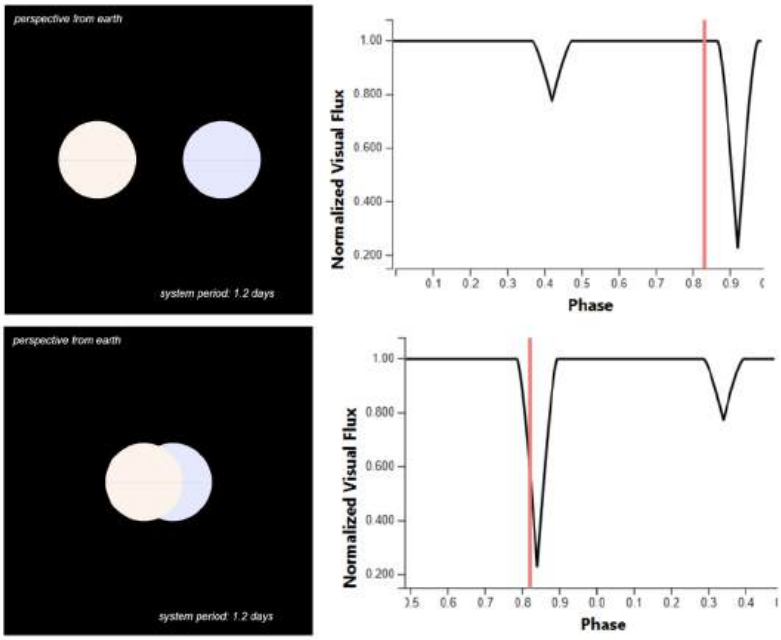


Fig 2 | A rudimentary simulation depicting eclipsing binaries and their resultant flux information as detected by a satellite (Eclipsing Binary Simulator).

As depicted in the geometry of Columbia University's Center for New Media Teaching and Learning (CCNMTL) Eclipsing Binary Simulator, a dip in the observed flux is detected upon the overlapping of these two stars (as shown by the red line, depicting the simulation’s current phase), with a greater dip being seen when the hotter star is eclipsed by the colder star rather than in the opposite scenario. This occurs as hotter stars radiate much more energy across a greater range of wavelengths as per Wien’s Displacement Law, therefore causing a greater dip in observed flux when this energy is then blocked by the colder star. Therefore, the temperature difference between two eclipsing binaries can be observed by the extent to which observed flux has changed.

Such dips in the observed flux are quite likely to be mistaken for exoplanet detection given that both eclipsing binaries and exoplanets exhibit the same patterns in transit depth during an eclipse. This issue is generally addressed in two ways (Almenara and Deeg; Brown). Firstly, two successive transits may be taken: if the second transit does not correlate with the first one, there is a greater likelihood that it is an eclipsed binary or another exoplanet eclipsing the star at that time. Secondly, other methods such as radial velocity are utilised in order to substantiate the

initial finding of an exoplanet, acting as a strong method of validation as they are not suspect to the same drawbacks.

However, there is also a third method (Gutiérrez-Canales et al.) which forms the primary focus of this paper, which can arguably prevent false positives: star centroids.

Star Centroids

Star centroids are the sub-pixel position of the star's assumed point source of light. In other words, they are an estimation of the exact centre point of the star, which generally coincides with the brightest point of the star from the perspective of planet Earth (Gutiérrez-Canales et al.). Note, however, that due to noise, the brightest point of the star may not directly signify the centre of the sun when imaged using conventional photometry methods. Therefore, for the best sub-pixel approximation, the intensity-weighted center of the star is calculated by taking a mean average of the intensities, or 'weights' of all the pixels.

The exact methodology involved in calculating the intensity-weighted center is not mentioned in this study as the TESS Science Processing Operations Centre (SPOC) provides the star centroid coordinates directly for each and every timestamp.

Ideally, during the transit, these centroid coordinates should not fluctuate as there is only a single source of light: the target star that is being observed. However, in the case of eclipsing binaries, since there are two sources of light, the centroid begins to move from one star to the other, depending on the brightness of the stars themselves as well as the specific time frame of the transit. For example, consider star A and star B, where star A happens to eclipse star B during transit and star B is the one under observation. From the perspective of the viewer (on planet Earth), the centroid of star B will seem to 'drift' to the direction at which star A is beginning to eclipse star B, as star A's apparent brightness influences that of star B, causing the perceived 'brightest point' to seem to shift towards star A and away from star B. As star A nears the end of its eclipse, it will then exhibit the opposite effect, with centroid positioning seeming to move from star A back to star B. Of course, if star A was just a planet, none of this movement would occur as the planet does not emit any light whatsoever.

Convolutional Neural Networks

A Convolutional Neural Network (CNN) is a type of neural network commonly used for training spatial data with multiple dimensions. This makes it highly ideal for time-series data, where each of the input features are measured between equally sized cadences, thus giving the dataset a spatial aspect which the model can make predictions upon to a high accuracy. In fact, this spatiality is the reason why many state-of-the-art exoplanet vetting algorithms such as ExoMiner (Valizadegan et al.) and AstroNet use CNNs for classifying exoplanets against false positives, just like in this study.

This model fundamentally uses a 'convolutional' layer, which applies filters on to the spatial data provided in order to extract the most important features present through convolution.

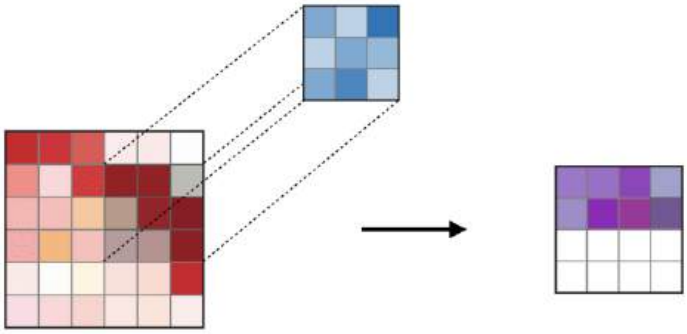


Fig 3 | A visual representation of the convolutional layer
(CS 230 - Convolutional Neural Networks Cheatsheet)

This is often followed by a second 'pooling' layer, which downsamples the data present by either max pooling or average pooling: max pooling is when the highest value in the given filter is taken; average pooling is when an average of all values in the given filter is taken. It is these values that are then passed onto the following layer.

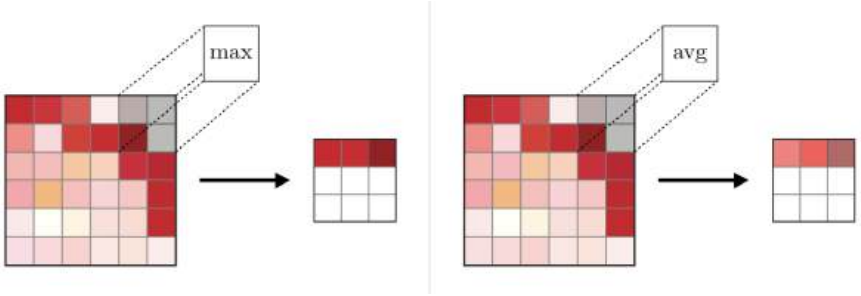


Fig 4 | A visual representation of the pooling layer
(CS 230 - Convolutional Neural Networks Cheatsheet)

Of course, additional layers such as can also be present in order to improve the inference results of a CNN, however they can vary from model to model depending on the context and the goal that the model is supposed to achieve.

A dual-branch CNN (Cui et al.), which will be a part of this study, is a variant of a CNN that has a 'local' branch and a 'global' branch. Each branch contains its own respective layers with their own respective parameters. The local branch aims to further scrutinise the data, getting a close look at finer details present within the data, whereas the global branch aims to get a generalised idea of the data and its broader patterns.

Research Question

It has been often underscored that the perceived change in centroid positioning, or the centroid movement, of a star can act as a viable method to differentiate exoplanetary detections from false positives such as eclipsing binaries. As discussed in the introductory section of this

study, whereas centroid movement should be near-zero when an exoplanet is transiting its star, it will ideally increase when a star eclipses another, since the perceived 'brightest point' calculated by the weighted average of all flux data should appear to move as the second star's flux influences that of the first.

However, this does beg an important question: "To what extent does centroid movement contribute in the vetting of exoplanets against eclipsing binaries?"

2 | Methodology

Fetching The Training Data

Gathering Examples from TOI Catalog

Before the input features can be fetched, the examples themselves need to be identified along with their corresponding output labels. In order to do so, the MAST Tess Objects of Interest (TOI) catalog has been selected as it contains several thousands of examples that have been already classified through alternative means of verification (such as radial imaging). Moreover, it contains each example's *TIC ID*, which is a unique identifier that is given to all observed celestial bodies observed by TESS during its operation. It is this identifier that will be used in order to download the time-series input features directly from SPOC (Jenkins et al.).

For this study, a total of **2,565 examples** have been fetched from the TOI, of which 1,324 examples are True Positives (TPs) and 1,241 examples are False Positives (FPs). The labels 'KP', 'PC' and 'CP' have been chosen for the true positives, whereas 'FP', 'EB' and 'NEB' have been chosen for the false positives. This provides a class imbalance of approximately 1:1.067, which lies within reasonable bounds thus class balancing has not been implemented. All examples have been stored in a manifest *csv* file which can then be referred to when the training data is being downloaded.

Gathering Input Features

After the *csv* file has been generated, a separate script reads the manifest and gathers the *TIC ID* of each example, and then proceeds to sequentially:

1. Query the MAST server for all the examples in the manifest that have been processed by SPOC (Jenkins et al.). The processing involves generating light curves from the discrete, raw data that is fetched from TESS's optics.
   Any examples whose relevant data contains too many missing values, is flagged by SPOC (Jenkins et al.) as unreliable, or contain no lightcurve data is discarded at this stage.
2. Download the TESS SPOC lightcurve files for each example, storing them in a temporary location for caching purposes.
3. Extract the relevant columns, most notably being *PDCSAP_FLUX, MOM_CENTR1* and *MOM_CENTR2. TIME* is also extracted for use during the feature engineering process (to commence phase folding), however it will not be included in the final training set.

4. Map the outputs to binary values ‘0’ and ‘1’ for the CNN model.
5. Store the contents - that being 1,000 time series values for each column - in a *csv* file along with the corresponding output labels for each example.

Note that only 2,218 of the original 2,565 have been successfully downloaded and saved in the curated dataset, as the rest of the 425 examples did not meet the required data quality standards or had missing values. It is worth noting that this has also affected the class imbalance from a 1:1.067 ratio to a 1:1.28 ratio, which is a minor imbalance which will not affect the model’s ability. Specifically, 942 false positive and 1246 confirmed exoplanets have been stored in the training set.

Feature Engineering

Before the model can be trained on the training set, the data itself must be processed in order to optimize it as much as possible, ensuring that the pattern that lies within the data is clear. To achieve this, the following was employed:

1. Sector normalisation, where the lightcurve data for all individual sectors of an example are concatenated into a single lightcurve using Lightkurve’s stitch()function. Each sector’s lightcurve is normalised to a common baseline prior to concatenation.
2. Sigma clipping, where any flux values deviating more than 5 standard deviations from the global median are discarded (VanderPlas and Ivezic). This is done in order to remove any anomalies in the data that the SPOC did not flag as unreliable.
3. Perform an out-of-transit normalisation (Ansdell et al.) that aims to identify all data points outside of the transit itself, since they are not required for the classification process. It essentially multiplies the transit period by ±1.5 in order to ensure that ingress/egress values are not discarded. This process is applied to both centroid and flux information, where the identified ‘out-of-transit’ data is normalised by dividing by the median and then standardised by subtracting that same median and dividing by the out-of-transit standard deviation.
4. Phase folding (Shallue and Vanderburg) where all cadences are mapped onto a singular normalised phase axis via:

$$x_0 = \frac{time - epoch}{period}$$

$$x = x_0 + \lfloor x_0 \rfloor \tag{4}$$

Where x refers to the folded phase. This is then shifted to $[-0.5, 0.5]$, with the transit itself being centered around phase 0 to center it. This is done to collapse the multitudes of observations into a single representation of the transit flux and centroid information, which is required for the fixed-length/shape input that the CNN will require.

5. Median binning (Shallue and Vanderburg) to 1,000 bins across $[-0.5, 0.5]$, producing an overall input (N, 1,000, 3) which will then be fed into the model.
6. Discarding any examples with more than 15% of those 1,000 bins being empty since interpolation (or attempting to fill in this empty data) would lead to too much synthetic information for the dataset.
7. Linearly interpolating examples with less than 15% of missing bin values via nearest neighbor, done by NumPy's built-in interp function on all 3 channels.
8. Finally, removing any outstanding values that are not in the range $[-10, 10]$ so that biasing does not occur due to these anomalies.

Training the Convolutional Neural Network (CNN)

For this study, a dual-branch 1-Dimensional CNN (Cui et al.) has been utilised via Tensorflow (Abadi et al.), with the input shape of the training data being (2003, 1000, 3) where there are 2,000 examples, 3 features per example, and 1,000 timestamps per feature. This model has been selected as it is known to perform particularly well with data represented as a time-series, as is the case with 1,000 timestamps spaced equally in time. Specifically, it leverages a local and a global branch, where the local branch is more scrutinous and sensitive to details present in the dataset since it only receives 200 of the 1,000 bins, whereas the global branch gives a general overview on the patterns found in the data (such as the transit variability).

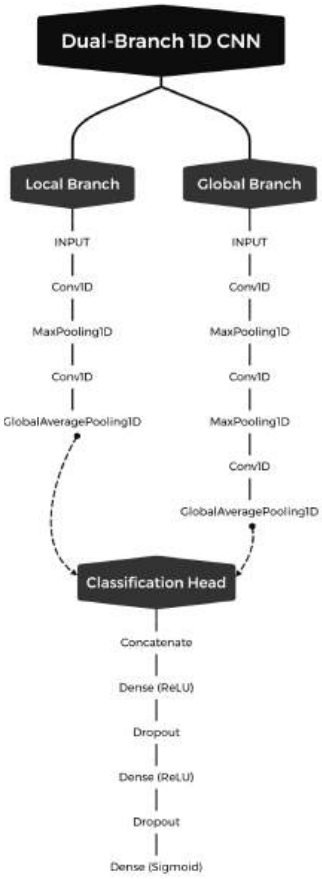


Fig 5 | A flowchart depicting an overview of the layers present in the neural network. Created using Xmind.

Parameters

For the model, a FocalLoss function was used rather than BinaryCrossEntropy ("CS229_lecture_notes.Pdf"), mainly in order to focus gradient boosting at decision boundary data rather than data that the model found easy to classify, which is a feature that the latter does not have. Specifically, the value $\gamma = 2.0$ was utilised so that predictions with a high confidence contributed less to the loss weights than lower-confidence predictions. A label smoothing of 0.1 was also used in order to regularize outputs and prevent saturated values.

Moreover, Adam ("CS229_lecture_notes.Pdf") was used in order to optimize the learning rate, with a 'warmup' function where the learning rate slowly ramped up from 1×10^4 to 1×10^3 over the course of five epochs. This not only prevents excessive 'overshooting' during gradient descent, but also allows Adam to use

larger learning rates only when it has collected enough statistical data in order to make more stable updates to the gradient descent function.

Additionally, an EarlyStopping function with a patience of 15 epochs has been implemented in order to prevent the model from getting too complex and overfit everything.

Given the size of the dataset, data augmentation has also been added in order to increase the robustness of the model, given that it simulates the following: random noise through a gaussian function, variations in the light curve data itself by moving data around ±20 bins, and variations in transit depth via flux scaling. All 3 algorithms aim to make the most of the limited number of examples present for the model to use while ensuring minimal synthesis of 'fake' data.

Note that **3 copies** of the model above have been created, but have distinctions that are as follows:

- Model A: Contains only flux information (2003, 1000, 1)
- Model B: Contains centroid and flux information as inputs (2003, 1000, 3)
- Model C: Contains only centroid information (2003, 1000, 2)

This distinction plays a crucial role in the goal of this study given its ability to isolate the fundamental inputs that the model has access to during training, ceteris paribus. Specifically, Model A and B will be compared in order to understand whether centroid movement information provides any additional support to the model when vetting exoplanets compared to flux alone. Model C merely acts as a control to ensure that centroid movement does indeed have a signal that the model can identify and classify upon.

Evaluation Metrics

In order to evaluate the model, a Stratified K-Folds with 10 folds has been selected as a method of cross-validation given its ability to use the entire dataset for testing the model without any data leaks. Not only are the examples randomly split into the 10 folds, but they are split in a manner such that each fold contains a near-even number of each class, hence being 'stratified'. Per-fold threshold optimization ensures that the thresholds are calculated in order to maximize the F1-macro score, which in turn boosts accuracy.

Moreover, an ensemble-level model is also generated by taking the average vector probabilities of each fold, with accuracy metrics being tested here as well.

In terms of the metrics themselves, a confusion matrix is output, clearly showing exactly where each example is being classified with respect to the target label that the example is supposed to be classified as.

On top of this, the AUROC score (Tafvizi et al.), used to determine the strength of the pattern or signal present in the dataset that the model can discern, is utilised to gauge the extent to which the model can find a pattern and generalise based on the data. A value of 0.5 means that

there is no pattern and the model is effectively guessing 50/50, and a value of 1.0 means that the pattern is very clear and every model can be appropriately classified.

It is denoted by:

$$AUROC = \int_0^1 TPR(FPR)\, dFPR \tag{5}$$

The calculation essentially takes the area under the True Positive Rate (TPR) curve with respect to the False Positive Rate (FPR) across all classification thresholds, where:

$$TPR = \frac{TP}{TP + FN}, FPR = \frac{FP}{FP + TN} \tag{6,7}$$

The F1 macro score is also calculated in order to gauge how well the model is classifying examples while being neutral and non-biased to class imbalance. Whereas the model does not exhibit a major class imbalance, it is worth using over accuracy which can tend to be more misleading.

The model uses the following formula:

$$F1_{macro} = \frac{\sum_{i=1}^{N} F1_n}{N} \tag{8}$$

Where N is the total number of classes (in this case 2) and:

$$F1 = \frac{2 \times TP}{2 \times TP + FP + FN} \tag{9}$$

F1 weighted scores are also used in order to get an idea of the score of the model given that all classes are weighted as per their abundance in the dataset.

Finally, in order to cross-compare the effectiveness of the models, this study will utilise the Wilcoxon signed-rank test (Wilcoxon), providing a direct value which depicts the similarity between two models in terms of their respective performance. It is this metric that will determine whether centroid information provides a meaningful difference on the accuracy and effectiveness of exoplanet classification.

3 | Results

After running all 3 models, each model produced its own confusion matrix, as well as metrics pertaining to how well the model ran. In particular, it is the ROC-AUX and F1 (macro + weighted) scores that depicts the effectiveness of the models, as well as the Wilcoxon signed-rank test that allows for cross-comparisons between these metrics and determines the

extent to which model accuracy has changed due to the removal of centroid information. Uncertainty values were calculated using the standard deviation experienced by the model.

With that in mind, here are the results across the 10 folds:

Model	Type	AUROC	AUROC Uncertainty (±)	F1-macro	F1-macro Uncertainty
A	Flux	0.7563	0.0337	0.6882	0.0409
B	Flux + Centroid	0.7642	0.0341	0.6862	0.0268
C	Centroid	0.6091	0.0395	0.5936	0.0330

The Wilcoxon test also yielded significant results, with:

Models Compared	W value	P-value
A and B	14.0	0.1934
A and C	0.0	0.0020

One direct observation suggests that whereas the models are able to discern a pattern in the dataset, classification performance itself isn't at the level of frontier vetting models such as ExoMiner or AstroNet, as neither of the models seemed to surpass an AUROC of 80%, much less. Moreover, the F1-macro score across all 3 models is considerably poor, staying under 70%. It can be observed that model C in particular performed significantly worse than models A and B across both AUROC and F1-micro. The uncertainty values are significantly high as well, remaining at the equivalent of 3-4% for AUROC and 2-4% for F1-macro.

4 | Evaluation

This study was built around the concept of investigating the extent to which centroid movement affects the classification of exoplanets against eclipsing binaries through the method of ablation. In order to identify this extent, the results of models A and B are to be scrutinised the most, as the theoretical gain in vetting efficacy by model B with regards to model A depicts the additional benefit that centroid movement motion is presenting to the model. Model C itself acts as a control present in order to confirm that centroid motion is indeed providing useful information that the model can pick up on and classify against.

As can be seen in the results, the AUROC values of models A and B are 0.7563 and 0.7642 respectively, boasting a difference of 0.0079. Whereas model B - the model containing flux and centroid information - performed marginally better than model A, it is worth noting that this value falls well into the uncertainty of both models at 0.0337 and 0.0341 respectively. Therefore, it can be said that centroid movement information has seemingly no effect on classification performance for a 1DCNN machine learning model. This insignificance of centroid motion is further supported by the Wilcoxon tests, where a p-value of 0.1934 deviates

significantly from the standard significance threshold of 0.05, implying that there is statistically no difference in the performance between models A and B and any difference observed occurred due to chance (Habibzadeh). This is supported by the W-value of 14.0 between models A and B, a considerably large value implying that there is no statistically significant difference between the two models' performances.

However, it is worth noting that Model C, which only had access to centroid information, did still possess an AUROC score of 0.6091, which is noticeably greater than 0.5 which would otherwise imply random and chance-based classification. This could be attributed to feature dominance present in Model B where the model purposefully ignores centroid movement information in favor of flux as the latter has a more discernible pattern. This is supported by a W-value of 0.0 between models A and C, implying that every single one of the 10 K-folds favored flux over centroid information, supported by the p-value of 0.0020 which is the floor threshold of the p-value given that only 10 K-folds were taken (and $2^{-9} = 0.00195 \approx 0.002$) proving a genuine statistical difference between flux and centroid movement in the models' evaluations.

Discussion

The largest limiting factor present in this study relates to the training data itself: there are not enough examples available for training the model in order to achieve a high accuracy without the risk of severe overfitting. This is especially prominent given the coarse pixel scale of 21 arcseconds per pixel utilised by the TESS telescope itself, making it exceedingly difficult to differentiate (and thus classify) between valid information and noise present in data and can cause serious misclassification up to a point of overfitting.

Consider popular ML-assisted exoplanet vetting tools such as AstroNet or ExoMiner (Valizadegan et al.): not only did they train their model on at least 15,000 different examples - that being 7.5 times the number of examples in the model produced for this study - but they also utilised a large quantity of Kepler data in order to train their model before migrating it for use on both TESS and Kepler data through transfer learning. Not only are there more examples present to train the model, but Kepler's pixel scale of 3.98" implies the presence of cleaner and higher quality data with a very low signal-noise-ratio (SNR) rendering it ideal for training purposes. Even a close literature equivalent, Ansdell et al., leveraged 15,737 classified Kepler examples and only reported a 1.7-2.5% improvement when augmenting centroid information in their ExoNet deep learning model, making it a genuine possibility that this improvement was obfuscated in the coarser TESS data used in this study.

Moreover, when taking a look at the Wilcoxon test taken between Models A and C, it is clear that the p-value is 0.002, being right below the threshold of 0.05. Whereas it can be said that a higher statistical power will increase result accuracy as the p-value is at the floor of the current statistical power of n=10, the p-value will still remain under 0.05 which is a finding significant enough to prove the drastic difference between centroid and flux information during classification on a 1DCNN.

Future Work

In the future, the model could be improved through many factors, most notably the use of a larger dataset. Given that the TOI catalog used to gather pre-classified examples only had a total of 1,224 false positives, the model was severely restricted with a lack of examples to train upon. Not only would this provide more examples to train the model upon, but it would permit the model to be more complex, with a larger number of branched nodes significantly reducing the risk of overfitting while allowing the model to classify with a greater accuracy. Moreover, this ablation study could use Kepler-trained models to classify the dataset via transfer learning, significantly overcoming TESS's limited pixel scale of 21" as opposed to that of Kepler at 3.98" (Valizadegan et al.). Proven successful by existing literature, this can allow for a much greater AUROC and accuracy across all 3 model variations as the models would have access to higher-quality training data. Lastly, the model could be refined in order to reduce feature domination by flux intensity to force Model A to consider both input channels equally rather than disregarding centroid information over much clearer flux intensity data.

5 | Appendix

All code has been uploaded to a GitHub Repository (Sharma).

Note that portions of the code were developed with the assistance of Claude Code, wherein the author has guided the logic, implementation and framework of the code and Claude aided in the iterative code generation process as well as debugging. All code has been thoroughly scrutinized, corrected and tested for validity. *No AI or LLM was used in the writing of this manuscript.*

Works Cited

Abadi, Martın, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems.

Almenara, J. M., and H. J. Deeg. “Rate and Nature of False Positives in the CoRoT Exoplanet Search.” Astronomy & Astrophysics, vol. 506, no. 1, Oct. 2009, pp. 337–41. arXiv.org, https://doi.org/10.1051/0004-6361/200911926.

Ansdell, Megan, et al. “Scientific Domain Knowledge Improves Exoplanet Transit Classification with Deep Learning.” The Astrophysical Journal Letters, vol. 869, no. 1, Dec. 2018, p. L7. Institute of Physics, https://doi.org/10.3847/2041-8213/aaf23b.

Ar, Pranav. Refining Planetary Radius Calculations by Exploring the Geometry of Transit Photometry.

Brown, Timothy M. “Expected Detection and False Alarm Rates for Transiting Jovian Planets.” The Astrophysical Journal, vol. 593, no. 2, July 2003, p. L125. iopscience.iop.org, https://doi.org/10.1086/378310.

Charbonneau, David, et al. “Detection of an Extrasolar Planet Atmosphere*.” The Astrophysical Journal, vol. 568, no. 1, Mar. 2002, p. 377. iopscience.iop.org, https://doi.org/10.1086/338770.

CS 230 - Convolutional Neural Networks Cheatsheet. https://stanford.edu/~shervine/teaching/cs-230/cheatsheet-convolutional-neural-networks/. Accessed 27 Mar. 2026.

“CS229_lecture_notes.Pdf.” https://cs229.stanford.edu/main_notes.pdf. Accessed 7 Apr. 2026.

Cui, Zhicheng, et al. “Multi-Scale Convolutional Neural Networks for Time Series Classification.” arXiv:1603.06995, arXiv, 11 May 2016. arXiv.org, https://doi.org/10.48550/arXiv.1603.06995.

Eclipsing Binary Simulator. https://ccnmtl.github.io/astro-simulations/eclipsing-binary-simulator/. Accessed 27 Mar. 2026.

Gutiérrez-Canales, F., et al. “Detecting False Positives with PLATO Using Double-Aperture Photometry and Centroid Shifts.” Astronomy & Astrophysics, vol. 707, Mar. 2026, p. A2. arXiv.org, https://doi.org/10.1051/0004-6361/202453294.

Habibzadeh, Farrokh. “The P Value: What It Is and What It Is Not.” Journal of Korean Medical Science, vol. 40, no. 44, Nov. 2025, p. e321. PubMed Central, https://doi.org/10.3346/jkms.2025.40.e321.

Jenkins, Jon M., et al. “The TESS Science Processing Operations Center.” Edited by Gianluca Chiozzi and Juan C. Guzman, Edinburgh, United Kingdom, 2016, p. 99133E. DOI.org (Crossref), https://doi.org/10.1117/12.2233418.

Levrard, B., et al. “FALLING TRANSITING EXTRASOLAR GIANT PLANETS.” The Astrophysical Journal, vol. 692, no. 1, Jan. 2009, p. L9. Institute of Physics, https://doi.org/10.1088/0004-637X/692/1/L9.

McDonald, Catriona H., and Dimitri Veras. “Binary Asteroid Scattering around White Dwarfs.” Monthly Notices of the Royal Astronomical Society, vol. 520, no. 3, Apr. 2023, pp. 4009–22. Silverchair, https://doi.org/10.1093/mnras/stad382.

Oshagh, M. Noise Sources in Photometry and Radial Velocities. 2018, pp. 239–49. arXiv.org, https://doi.org/10.1007/978-3-319-59315-9_13.

Shallue, Christopher J., and Andrew Vanderburg. “Identifying Exoplanets with Deep Learning: A Five-Planet Resonant Chain around Kepler-80 and an Eighth Planet around Kepler-90.” The Astronomical Journal, vol. 155, no. 2, Jan. 2018, p. 94. Institute of Physics, https://doi.org/10.3847/1538-3881/aa9e09.

Sharma, Advik. Mrbeats09/Yet-Another-Exoplanet-Classifier. 4 Mar. 2026, Python. Version 1.00, 7 Mar. 2026. GitHub, https://github.com/mrbeats09/Yet-Another-Exoplanet-Classifier.

Tafvizi, Arya, et al. “Attributing AUC-ROC to Analyze Binary Classifier Performance.” arXiv.Org, 24 May 2022, https://arxiv.org/abs/2205.11781v1.

Valizadegan, Hamed, et al. “ExoMiner: A Highly Accurate and Explainable Deep Learning Classifier That Validates 301 New Exoplanets.” The Astrophysical Journal, vol. 926, no. 2, Feb. 2022, p. 120. Institute of Physics, https://doi.org/10.3847/1538-4357/ac4399.

VanderPlas, Jacob T., and Željko Ivezic´. “PERIODOGRAMS FOR MULTIBAND ASTRONOMICAL TIME SERIES.” The Astrophysical Journal, vol. 812, no. 1, Oct. 2015, p. 18. Institute of Physics, https://doi.org/10.1088/0004-637X/812/1/18.

Wilcoxon, Frank. “Individual Comparisons by Ranking Methods.” Biometrics Bulletin, vol. 1, no. 6, 1945, pp. 80–83. JSTOR, https://doi.org/10.2307/3001968.

A Brief Exploration of Black Holes in the Early Universe By Reuben Saunders

Abstract

With the development of the James Webb Space Telescope (JWST), a great deal of new information concerning the early universe has come to light. This has created a need for a reexamination of previous theories on the early universe, particularly the universe within the first billion years after the Big Bang. A particularly interesting finding that has been analysed is the appearance of supermassive black holes (SMBHs) that have more mass than would be expected if the previously-held view on black hole formation is accepted. The proportion of the mass of the supermassive black holes to the mass of the rest of the host galaxy suggests that these black holes were formed before their galaxy and formed their galaxy. The ramifications of this theory actually being proven could have wide-ranging implications, and would require a new look at previously thought information on black holes. These black holes would form differently and impact the universe around them differently than other black holes would. A large number of recent papers have examined this subject, and this paper aims to review the information that is already out there and to develop an evidence backed and coherent theory based on said data.

Introduction

In modern cosmology, the problem of disparate growth speeds of celestial objects has been studied and catalogued, and as of the writing of this paper, no definitive answer to the quandary has been established. The specific difficulty established arises when the mass of SMBHs in proportion to the mass of their host galaxies. In galaxies and supermassives that are found within the first few hundreds of millions of years after the big bang, large quantities of "overmassive black holes" or black holes that have masses greater relative to their host galaxies than was thought theoretically possible by previously held ideas (McClymont et al.). These findings appear to suggest direct collapse as the mechanism behind the creation of early universe supermassives. Direct collapse is the creation of black holes due to large pockets of gas rather than stellar collapse or accretion at the center of galaxies (Zhang, Liu, and Bromm). This also suggests that early galaxies were formed with heavy pockets of dark matter, because this would allow the black hole of a galaxy to grow at a much faster rate than the stars and stellar mass of a galaxy (McClymont et al.). This paper aims to explore these effects, with the question of: How did black holes in the early universe affect the growth of celestial objects, and how were these black holes affected by the composition of the universe around them? From this question, an inference will be derived from this question about how black holes affect the universe around them. This inference will be justified through an explanation of current scientific literature on the subject. A literature review will allow for the collection and analysis of a plethora of information on this research question. Following this introduction, section two will be a collection of different papers on this topic, each with multiple sentences of summary and analysis, section three will be a discussion section, where the different ideas in these papers are brought together to form a theory, and section four will be a conclusion.

Literature Review

The Growth of both Supermassive and Primordial Black Holes in the Early Universe

This literature review will begin with an exploration of how these early universe and highly redshifted (far away in space and time) black holes came to be and how they accreted matter in the early universe. The formation of these black holes is incredibly different from the theoretical formation of black holes today. These black holes are theorized to form from direct collapse, which is as it sounds: a direct collapse of pockets of gas rather than a collapse of a star or at the center of a galaxy. The first paper exploring this that will be reviewed is A Novel Formation Channel for Supermassive Black Hole Binaries in the Early Universe via Primordial Black Holes (Zhang, Liu, and Bromm). This paper posits that the supermassive black holes formed in the early universe due to primordial black holes (PBHs). In this paper, primordial black holes are massive black holes that were formed through the direct collapse of specifically metal-poor gas pockets. It studies how the formation of massive primordial black holes could help seed the creation of other black holes in the early universe. The paper simulates what would happen to the gas cloud around a primordial black hole, as well as simulating how it would cause direct collapse and the formation of secondary black holes. The simulation proves that primordial black holes create a larger dark matter halo and attract surrounding gas from the gas pocket. It also showed that it created thermal feedback, which delays the cooling and collapsing of gas. This shows the two ways that black holes can regulate the universe around them, and especially galaxy growth. Even in today's universe, these black holes cause processes to regulate galaxy evolution. The simulation also found that the primordial black holes created a runaway collapse effect, proving that primordial black holes are a viable theory as to the existence of so many overmassive black holes in the early universe.

The following paper on this topic is The growth efficiency of high-redshift black holes (Pacucci et al.). This paper describes the possible growth of these early universe black holes to explain the results and observations that we see today. In the paper, there is a numerical simulation of how an early universe black hole seed at the center of a dark matter halo would grow, the same exact conditions as outlined in the first paper. It simulated both a high density and a low density black hole, allowing for a comparison between the results of the two. The paper found that for a high and low density black hole, the low density black hole will hit its Eddington Limit first, meaning that their growth will be stunted. It also found that low density black holes are more likely than high density black holes to fall into a flow where there is more ejection than accretion (feedback limited). The higher density seed was more likely to have more accretion than ejection (feedback dominated). This could theoretically cause the low density black holes to evaporate due to hawking radiation, while the high density black holes grow. This would provide a reason for the fact that we only see supermassive black holes in the early universe, and not any smaller black holes.

How Primordial and Early Supermassives Affected the Universe Around Them

Now that there is an established theoretical framework for how these early universe and primordial black holes came to be, there will now be an exploration of how these black holes impact the universe around them. How do Massive Primordial Black Holes Impact the Formation of the First Stars and Galaxies? (Zhang, Liu, Bromm, et al.). is a paper with a hydrodynamical simulation of a massive primordial black hole's impact on the formation of the first galaxy. The simulation found that primordial black holes regulate the growth of the first (population III) stars in the same way that it regulates black hole growth. The seeding of collapse speeds up the development and feedback slows down the development of these stars. The simulation also found that the primordial black hole created a more favorable environment for the development of the first stars than the environment that would have existed if the primordial black holes did not exist. The simulation also found that primordial black holes create galaxies that are dense star clusters. Some JWST data that has been found of these early universe galaxies is consistent with this, providing observational evidence for this idea. Without the existence of these primordial black holes, our universe would have celestial objects nearly as prevalent or massive. These black holes are a key factor in the evolution of our universe.

A separate simulation is found in The formation of the first quasars: the black hole seeds, accretion, and feedback models (*Formation of the First Quasars: The Black Hole Seeds, Accretion, and Feedback Models | Monthly Notices of the Royal Astronomical Society | Oxford Academic*). This article includes a simulation of a quasar, which is a high energy luminous object created by the accretion surrounding a black hole. This quasar influenced the surrounding area so much that it created stars, much in the same way as in the other article. The most interesting finding from this simulation though was when the stars started forming. The simulation found that stars would not form until a black hole grew to a certain mass. This would allow for a black hole to have time to accrete before the surrounding galaxies even began to grow. This would be a very important consideration when trying to evaluate why there are so many overmassive black holes. This is another possible reason as to why we see so many overmassive supermassives in the early universe. These 4 pieces of literature all suggest how black holes regulated the early universe. These black holes shaped the world that we live in today.

Discussion

If these theoretical ideas about black holes in the early universe are conclusively proven, the implications for physics would be massive. Previously, the belief about black holes was that they either formed from the collapse of a dying star or formed at the center of an already established galaxy. If it is proved that there are black holes that collapsed and then formed a galaxy around themselves, it would completely rewrite the order of how these processes are thought to happen. This would change how we think about the universe, especially with some ideas about primordial black holes as dark matter. Recent information from LIGO has found that primordial black holes may constitute an important fraction of dark matter in the universe (Clesse and García-Bellido). These primordial black holes would answer one of the questions

that has plagued physicists for years, which is that there is a great chunk of matter that should exist but has not been observed. Primordial black holes could be the answer to two of the most important questions that are unanswered in physics. Those are of what dark matter is and how the first stars and galaxies in our universe came to be. From this one concept backed up by observations from the James Webb Space Telescope. All four of these papers point to the same collective hypothesis. They have established a picture of the early universe in which primordial black holes are formed by direct collapse, leading to the formation of other black holes and stars that grow faster than their galaxies, creating overmassive black holes. All of them appear to agree with this picture, and this collection can be used to be built upon for future work.

Conclusion

This paper answers the question of how primordial and early supermassive black holes interacted with the environment around them. The current literature shows that black holes were a main regulator of the universe around them, much in the same way that black holes act today. This project has taught me on how black holes regulate the environment. Through this I have inferred that black holes can both seed and repress growth, and that black holes can on their own create the systems that we see in our universe today. This paper has gathered research on early universe black holes to create a coherent theory on how the early universe worked with and around them. From this baseline idea, more focused observations can be made. With an established picture of what these black holes look like, we can now try to find specific examples of these already established examples for primordial black holes. The next step is to build even more powerful telescopes, so that we can find even earlier universe examples of the first primordial black holes. One of the telescopes that should be built is the Lunar Crater Radio Telescope. The Lunar Crater Radio Telescope is a radio telescope made from a mesh wire in a crater on the far side of the moon. This telescope will allow us to do research on the cosmic dark ages, which was the time before the first celestial objects came into existence.(Goel et al.) Another telescope that will be useful is the Laser Interferometer Space Antenna, or LISA. LISA will allow us to test that theory of primordial black holes as dark matter by analysing the gravitational waves.(*Testing Primordial Black Holes as Dark Matter with LISA | Phys. Rev. D*) More funding for this research and the building of these telescopes will allow us as a species to understand more about the universe we live in, and is an important step in advancing for a better future.

Works Cited

Clesse, Sebastien, and Juan García-Bellido. “Seven Hints for Primordial Black Hole Dark Matter.” arXiv:1711.10458, arXiv, 29 May 2018. arXiv.org, https://doi.org/10.48550/arXiv.1711.10458.

Formation of the First Quasars: The Black Hole Seeds, Accretion, and Feedback Models | Monthly Notices of the Royal Astronomical Society | Oxford Academic. https://academic.oup.com/mnras/article/514/4/5583/6604892. Accessed 17 July 2026.

Goel, Ashish, et al. “Probing the Cosmic Dark Ages with the Lunar Crater Radio Telescope.” arXiv:2205.05745, arXiv, 2 June 2022. arXiv.org, https://doi.org/10.48550/arXiv.2205.05745.

McClymont, William, et al. “Overmassive Black Holes in the Early Universe Can Be Explained by Gas-Rich, Dark Matter-Dominated Galaxies.” Monthly Notices of the Royal Astronomical Society, vol. 545, no. 1, Jan. 2026, p. staf2092. Silverchair, https://doi.org/10.1093/mnras/staf2092.

Pacucci, Fabio, et al. “The Growth Efficiency of High-Redshift Black Holes.” Monthly Notices of the Royal Astronomical Society, vol. 452, no. 2, Sept. 2015, pp. 1922–33. DOI.org (Crossref), https://doi.org/10.1093/mnras/stv1465.

Testing Primordial Black Holes as Dark Matter with LISA | Phys. Rev. D. https://journals.aps.org/prd/abstract/10.1103/PhysRevD.99.103521. Accessed 21 July 2026.

Zhang, Saiyang, Boyuan Liu, and Volker Bromm. “A Novel Formation Channel for Supermassive Black Hole Binaries in the Early Universe via Primordial Black Holes.” The Astrophysical Journal, vol. 992, no. 1, Oct. 2025, p. 136. Institute of Physics, https://doi.org/10.3847/1538-4357/ae061c.

———. “How Do Massive Primordial Black Holes Impact the Formation of the First Stars and Galaxies?” The Astrophysical Journal, vol. 987, no. 2, July 2025, p. 185. Institute of Physics, https://doi.org/10.3847/1538-4357/adddb4.

How Does the Kauffman Bracket with its Writhe Correction Prove that the Jones Polynomial is a Well-defined Knot Invariant? By Zhiheng Luo

0. Abstract

This paper proves that the Jones polynomial is a well-defined knot invariant by constructing it from the Kauffman bracket and the writhe. We show that the Kauffman bracket is unambiguous, since its state-sum formula is independent of the order in which crossings are resolved, and that it survives Reidemeister moves II and III but fails move I. To repair this, we define a normalized invariant that multiplies the bracket by a compensating power of the writhe, and prove that this correction restores invariance under move I without disturbing invariance under moves II and III. A substitution then transforms this invariant into the Jones polynomial, which inherits both its invariance and its unambiguity. We close with a worked example computing the Jones polynomial of the trefoil knot.

1. Introduction

At its core, knot theory seeks to answer a deceptively simple question: How can we determine when 2 knots are identically the same? To solve this problem, mathematicians rely on tools known as “knot invariants". One of the most well-recognized and powerful invariants in the field is the Jones polynomial (Jones). This paper will explore the following question: "How does the Kauffman bracket serve as proof for the knot invariance and unambiguity of the Jones polynomial?"

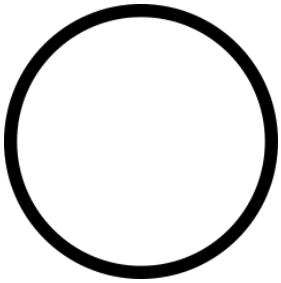
To thoroughly investigate this, the paper will proceed as follows: Section 2 will establish definitions and terms in knot theory such as Reidemeister moves. Section 3 will explain what the Kauffman bracket is and its utilities. Section 4 will test the Kauffman Bracket on Reidemeister Moves. Finally, Section 5 will introduce the Jones Polynomial and explore the proof of the polynomial’s invariance and unambiguity via the Kauffman Bracket.

2. Preliminaries

2.1 Definition of a knot and knot diagrams

In knot theory, a knot is defined as an embedding of a circle into three-dimensional Euclidean space R^3 (Adams). In other words, a knot is a closed loop that contains no loose ends and does not intersect with itself. Because three-dimensional (3D) knots are rather arduous to visualize, mathematicians study knots by projecting them onto 2D space. These 2D projections are known as “knot diagrams”. The simplest possible knot diagram is a circle, otherwise known as the “unknot”.

Fig 1. Diagram of an unknot



When a knot is represented as a diagram, the loop inevitably overlaps itself, creating intersections known as crossings. At every crossing, one segment of the knot passes over the other. To indicate this in a diagram, the lower line is broken, distinguishing the over-strand from the under-strand.

Fig 2. Diagram of a crossing

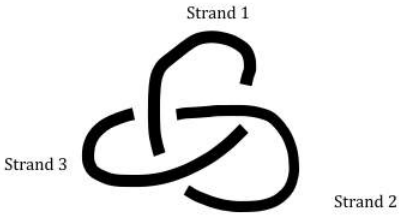


As seen in Figure 2, a single crossing is composed of one over-strand (the longer, unbroken line) and two under-strands (the shorter, broken lines).

With the definition of a crossing established, we can define another part of the diagram: strands. A strand is a continuous segment of the knot that begins immediately after one undercrossing and ends at the next undercrossing it reaches. A strand may pass over any number of crossings along the way (passing over does not end it) but it terminates the moment it passes under another strand.

To gain a better understanding of how crossings and strands interact, we can analyze the trefoil knot.

Fig 3. Trefoil knot



As we can observe from the figure, the trefoil diagram contains three crossings. Consequently, it is composed of three strands. Each strand travels over one crossing before terminating at the undercrossings on both ends.

2.2 Reidemeister Moves

To explore the main question of knot theory (knot equivalence), there exists a test to prove if two diagrams represent the same knot. In 1927, mathematician Kurt Reidemeister proved that any deformation of a knot can be broken down into a sequence of three maneuvers (Roberts), which are known as the Reidemeister moves.

Before addressing all the maneuvers, it is essential to establish the rule of continuous deformation, formally referred to as planar isotopy (or Reidemeister move 0). This equivalence allows for the stretching, bending, or shrinking of any segment of a strand within the diagram, provided that the strand never crosses another strand. Under planar isotopy, the identity of the diagram remains unchanged. Note that this step will be trivially used to clarify some diagrams.

Fig 5. Before and after Reidemeister move 0/planar isotopy



Once we account for planar isotopy, all other transformations can be done by Reidemeister moves:

- Reidemeister move I (Twist):
  This move takes a strand of the knot and twists it into a loop or curl, which creates a new crossing. Conversely, this move allows you to untwist an existing curl to straighten the strand. This move introduces or removes a crossing in a knot, altering the diagram's total crossing and strand number by one.

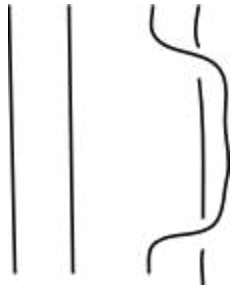
Fig 6. Before and after Reidemeister move I



- Reidemeister move II:

This move involves taking one strand and putting it under (or over) an adjacent strand. Because the strand enters and exits the crossing area, this action creates two new crossings simultaneously. Similar to Move I, this can be done in reverse: A strand that is laying over another can be pulled away, ridding of two crossings at once.

Fig 7. Before and After Reidemeister move II



- Reidemeister move III:
  Unlike the first two moves, Reidemeister Move III does not change the total number of crossings in a diagram but rather rearranges them. This move involves taking a strand and slides it past an existing crossing made by two other strands.

Fig 8. Before and after Reidemeister move III



These moves serve as a benchmark in knot theory. If two diagrams can be connected by a sequence of Reidemeister moves (including Move 0), they are proven to be the same knot. Furthermore, for a polynomial to be proven as a knot invariant, it must survive all three of these crossing maneuvers without changing its value.

2.3 Knot Invariant

Because testing Reidemeister moves by hand to determine if two diagrams represent the same structure is virtually impossible or requires a plethora of computational power for complex knots, we will utilize a tool known as knot invariants. A knot invariant is a property: A number, a group, or a polynomial that remains unchanged regardless of how much a knot is deformed under the Reidemeister moves.

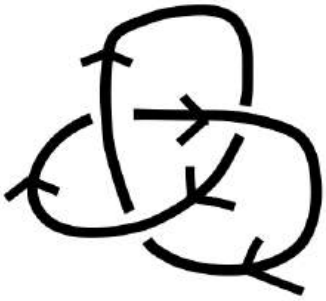
Therefore, the rule for proving an invariant is as follows: For a polynomial to be proven as a knot invariant, it must be demonstrated that its value survives all three Reidemeister moves

without changing. If a calculation yields the same polynomial for two distinct knot diagrams, it provides strong evidence that the two diagrams may be the same knot. However, this isn't always true due to polynomials having blind spots.
This paper will focus on polynomial invariants, specifically the Jones polynomial. In the following sections, we will establish the base of the proof of their knot invariance.

2.4 Orientation and Crossing Signs

Before we can calculate knot polynomials, we first need to give the knot a direction. By default, a standard knot diagram is "unoriented," meaning it is directionless. To create an oriented knot, we must choose a direction of travel along the loop. In a 2D diagram, this is indicated by drawing arrows along the strands to track the path.

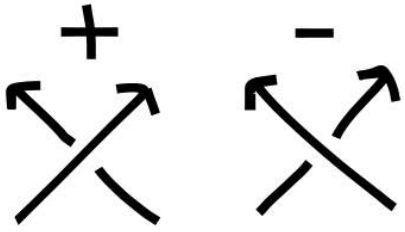
Fig 9. Trefoil knot with orientation



Once the knot has a direction, we can assign a value of either +1 or -1 to every single crossing in the diagram. This value is determined by looking at the crossing and observing the direction that the two intersecting strands are traveling.

A crossing's sign is determined by the following unambiguous rule: let O be the direction the over-strand travels away from the crossing, and let U be the direction the under-strand travels away from the crossing. Rotate O by 90 degrees counterclockwise. If the rotated direction points the same way as U, the crossing is positive (+1). If instead rotating O by 90 degrees clockwise aligns it with U, the crossing is negative (-1). Figure 10 below shows both cases explicitly, with the over-strand drawn unbroken and the under-strand drawn with a gap, so that which strand is on top is never ambiguous.

Fig 10. Positive and negative crossings



Orientation and crossing signs allow us to calculate a property known as the writhe, which will become essential later in this paper.

3. Kauffman Bracket

3.1 Foundations

Having established Reidemeister moves, we now introduce the primary topic of this paper: The Kauffman bracket. Discovered by Louis Kauffman in 1987, the bracket polynomial (Kauffman), denoted as $\langle K \rangle$, is a model that assigns a polynomial in variables A and A^{-1} to an unoriented knot or link diagram.

The bracket evaluates a diagram by recursively "smoothing" (cutting and reconnecting) every crossing until the knot is simplified into a collection of trivial, non-intersecting circles. The bracket follows three rules:

- Rule 1 (Trivial Knot): If ◯ represents the standard unknot, $\langle \bigcirc \rangle = 1$
- Rule 2 (Smoothing Relation): Every crossing can be resolved in two distinct ways. Standing at the crossing and facing along the over-strand, rotate the over-strand counterclockwise: the two regions it sweeps through before reaching the under-strand are called the A-regions, and the A-smoothing is the resolution that merges these two regions into one. The remaining two regions, swept clockwise, are the A^{-1}-regions, merged by the A^{-1}-smoothing instead. This labeling is only a convention: consistently swapping which smoothing is called A and which is called A^{-1} throughout the whole paper just produces a mirror-image, so it does not matter which one is chosen, only that the same choice is used at every crossing. The bracket is defined as the sum of these two states:

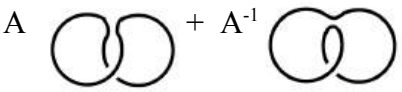


- Rule 3 (Disjoint Circle): If a diagram consists of a knot K alongside a disjoint, unknotted circle ◯, the trivial circle can be removed by multiplying the remainder of the bracket by $-A^2 - A^{-2}$:
$\langle K \sqcup \bigcirc \rangle = (-A^2 - A^{-2}) \langle K \rangle$

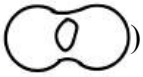
Example 3.1 The bracket polynomial of the Hopf Link can be calculated by applying these rules. We start with the Hopf Link:



Since the Hopf link consists of 2 crossings, we will have to choose which to resolve first. Note that in the following section we will prove that the crossing order does not matter.
Let's start with the top crossing. Following Rule 2, we get:



Now, since there is still a crossing unresolved, we apply Rule 2 again. We need to apply Rule 2 on both of these diagrams. We get:

A(A  + A^{-1} ) + A^{-1}(A  + A^{-1} )

We see the second and third diagram are both unknots, which Rule 1 states will equate to 1. Then, we will apply Rule 3 on the first and last diagram, both of which have 2 disjoint circles, this rule will multiply the diagrams by $(-A^2 - A^{-2})$. This gives us:

$$(A)(A)(-A^2 - A^{-2}) + (A)(A^{-1})(1) + (A^{-1})(A)(1) + (A^{-1})(A^{-1})(-A^2 - A^{-2})$$

Simplifying this, the equation cancels out to:

$$-A^4 - A^{-4}$$

Thus, we have calculated the bracket polynomial for the Hopf Link.

3.2 State-Sum Formula

A knot polynomial is well-defined if its calculation is independent of the "order of operations" (or in this case, the smoothing order). This section will prove that the Kauffman bracket is well-defined.

Let K be an unoriented knot diagram containing n crossings. We define a "state," denoted as S, as a configuration in which every single crossing in the diagram has been resolved into either an A smoothing or an A^{-1} smoothing.

By calculating all the possible states of the bracket, we can evaluate all crossings in the diagram simultaneously. This allows us to express the Kauffman bracket as a summation. Because each of the n crossings presents two choices, the knot diagram generates a set of 2^n possible states.

For every state S consists of a finite number of disjoint circles. To calculate the polynomial of any specific state S to the overall polynomial, we will define three parameters:

- $\alpha(S)$: The total number of crossings in state S resolved using the A smoothing.
- $\beta(S)$: The total number of crossings in state S resolved using the A^{-1} smoothing.
- Note that $\alpha(S) + \beta(S) = n$
- $|S|$: The total number of disjoint circles formed in state S.

According to the rules of the bracket, each A smoothing contributes a factor of A, and each A^{-1} smoothing contributes a factor of A^{-1}. Multiplying these coefficients across the entire state yields the term $A^{\alpha(S)-\beta(S)}$.

Furthermore, the rules dictate that every disjoint loop multiplies the term by $-(A^2 + A^{-2})$. Because one loop acts as the base, this factor is applied $|S| - 1$ times.
By combining these parameters and summing all 2^n possible states, the Kauffman bracket is defined by the state-sum formula:

$$\langle K \rangle = \sum_S A^{\alpha(S)-\beta(S)}(-A^2 - A^{-2})^{|S|-1}$$

By evaluating the diagram as a summation, the state-sum formula removes the choices of smoothing order and guarantees that the Kauffman bracket is a well-defined polynomial.

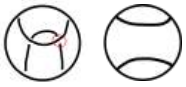
4. Proving Reidemeister Moves

For the Kauffman bracket to be considered a knot invariant, it must survive Reidemeister Moves I, II and III without its polynomial value changing. We will investigate each move, starting with R2.

4.1 Reidemeister Move II

We must prove that the before and after values of the Kauffman bracket for localized strands remain the same. To do this we can apply the Kauffman bracket onto the “after” of a Reidemeister move II.

After and before Reidemeister move II:



Starting with the “after” state, we will choose the right crossing (again, we proved that the smoothing order does not matter). By applying Rule 2, we get:

A  + A^{-1} 

Which turns into:

A(A  + A^{-1} ) + A^{-1} (A  + A^{-1} )

This simplifies to:

A^2  +  +  + A^{-2} 

Now, we can set these local diagrams as variables:
Let V represent diagrams with 2 vertical strands
Let H represent diagrams with 2 horizontal strands.

We must show that ⟨after⟩ = ⟨before⟩. Resolving the two crossings of the "after" diagram in sequence and tracking which of the four resulting states are planar-isotopic to which named diagram (V = two vertical non-crossing strands, H = the "before" diagram) makes this a direct verification. Note that the disjointed circle of  can be eliminated using Rule 3 to get:

$(-A^2 - A^{-2})V$

In total:

$A^2V + (-A^2 - A^{-2})V + H + A^{-2}V$.

We can see that the V variable cancels out completely, leaving us with only H.
Thus, the “after” state simplifies into the “before” state, showing that the Kauffman Bracket is unchanged by Reidemeister move II.

4.2 Reidemeister move III

We must prove that ⟨before⟩ = ⟨after⟩ for a Reidemeister move III. Both diagrams contain three crossings arranged identically except for the position of one strand relative to the crossing of the other two.

We will apply the Kauffman Bracket to the "after" state. Via Rule 2, we resolve the crossing involving the moving strand:

A  + A^{-1} 

The second diagram in this expansion still contains the two crossings not involving the moving strand, arranged as a Reidemeister move II configuration. Applying Reidemeister move II (which we have already proven leaves the Kauffman bracket unchanged) simplifies it to a fixed diagram we call P. The first diagram, which already has no crossing between those two strands, is a fixed diagram we call Q. This gives:

Q:  P: 

$A\langle Q\rangle + A^{-1}\langle P\rangle$

Now, we do the same thing to the “before” state: resolving the corresponding crossing using Rule 2, then apply Reidemeister move II to the remaining pair of crossings. This produces the same two diagrams, P and Q, with the same coefficients:

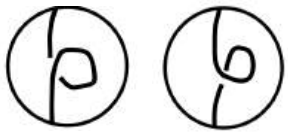
A  + A^{-1} 

Into $A\langle Q\rangle + A^{-1}\langle P\rangle$

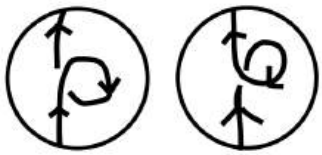
Because both the "before" and "after" states reduce to the same expression, independent of which side we started from, we conclude ⟨before⟩ = ⟨after⟩, proving that the Kauffman Bracket is unchanged by Reidemeister move III.

4.3 Reidemeister move I

Moving onto Reidemeister move I. However, in this case, we have two versions of the move that will lead to different results.
These are 2 possible "after" states of R1:



In this case, we will need to give each diagram an orientation to show the difference.

 signs, we can distinguish that the left diagram is a positive crossing, while the right one is negative.
Applying the Kauffman Bracket to each diagram, for the left diagram we get:

A  + A^{-1} 

For the right diagram, we get:

A  + A^{-1} 

Note that the "before" state of the move is , which we can substitute as a variable S.
We see that for the diagram with the positive crossing, there is a disjoint circle in the A coefficient. By applying Rule 3, we get:

$$A(-A^2 - A^{-2})S + A^{-1}S = -A^3S$$

On the contrary, if we apply Rule 3 on the diagram with the negative crossing, we get:

$$AS + A^{-1}(-A^2 - A^{-2})S = -A^{-3}S$$

Compared to the "before" state of S. We can see that a positive crossing alters the polynomial by multiplying it by $-A^3$, while a negative crossing alters it by $-A^{-3}$.

Since $\langle D^+\rangle = -A^3\langle D\rangle \neq \langle D\rangle, \langle D^-\rangle = -A^{-3}\langle D\rangle \neq \langle D\rangle$ The Kauffman Bracket is not a knot invariant under Reidemeister move I, even though it is unchanged by Reidemeister moves II and III. However, Section 5 defines a normalized invariant that repairs this defect.

5. The Jones Polynomial

5.1 The Writhe

Proposition 4.3 shows that the only obstruction to the Kauffman bracket being a knot invariant is its behavior under Reidemeister move I, where it is rescaled by $-A^{\pm 3}$ according to the sign of the crossing introduced. Since this rescaling is governed by crossing signs, we can cancel it by defining a new, normalized invariant: for an oriented diagram D, let $X(K) = (-A^3)^{-w(D)}\langle D\rangle$, where $w(D)$ is the writhe of D. We now prove that $X(K)$ is a knot invariant. The writhe, or $w(D)$, is the sum of all positive and negative crossings of a knot diagram. The writhe finds the total number of positive crossings and subtracts it by the total number of negative crossings. Since we know that positive crossings contribute a multiplier of $-A^3$ while negative crossings contribute a factor of $-A^{-3}$, writhe correction is able to counter the changes by the crossings by multiplying the polynomial by a factor of $(-A^3)^{-w(D)}$.

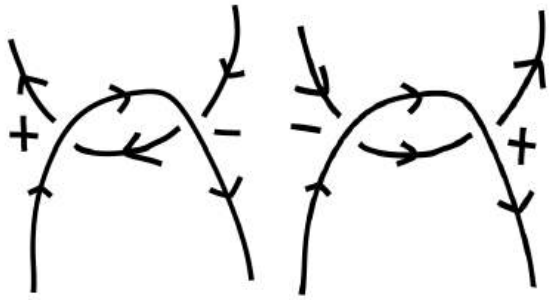
Proof: The total multiplier by positive and negative crossings could be written as this equation, where P is the number of positive crossings and N is negative crossings:
$(-A^3)^P * (-A^{-3})^N = (-A^3)^{P-N}$
Since $P - N$ is also known as the writhe (the writhe is defined by positive crossings minus negative crossings) it simplifies to:
$(-A^3)^{w(D)}$

We can see that the Kauffman Bracket is multiplied by this factor due to Reidemeister 1, which the writhe correction factor cancels out completely, leaving us with the regular Bracket polynomial. This effectively eliminates the problem of Reidemeister move I.

However, with the new writhe, we also introduce the problem of surviving Reidemeister moves II and III again.

Reidemeister move II

We have already proved that the Kauffman Bracket is unaffected by these moves, but we now need to prove that the written correction is also not altered by these moves.
To prove this, we need to assign orientation to the moves.

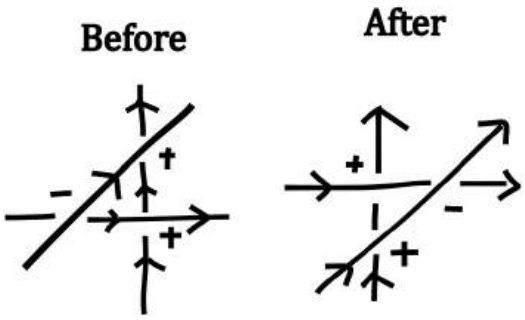


As we can observe, whatever direction the two strands are, one crossing will always evaluate to +1 and the other will always evaluate to -1. Because (+1) + (-1) = 0, applying R2 will

always add a net value of zero to the writer. Therefore, the total writhe of the knot remains unchanged and thus the corrected Kauffman Bracket is unaffected by R2.

Reidemeister move III

R3 is more straightforward.



Unlike R2, no new crossings are added. More importantly, sliding the strand does not flip any crossing signs. Because the intersection remains identical before and after the slide, the sign of each crossing stays the same. With no signs flipping, the sum remains constant, meaning the net change to the writhe under R3 is zero.

Ultimately, because the writhe is immune to R2 and R3, we can pair it with the Kauffman bracket (which is also immune to R2 and R3) to build a knot invariant.

Thus, the polynomial $(-A^3)^{-w(D)}\langle K\rangle$ is an invariant under all three Reidemeister Moves.

5.2 Invariance of Jones Polynomial

Having corrected the Kauffman bracket, which we can express as $X(K)=(-A^3)^{-w(D)}\langle K\rangle$, we can actually find a direct relationship with the Jones Polynomial. To change the corrected Kauffman bracket to the Jones polynomial, denoted as $V_K(t)$, we can apply the substitution $A = t^{-1/4}$ (Since A is arbitrary). We get:

$$V_K(t) = X(K)\Big|_{A=t^{-1/4}}$$

Because $X(K)$ has already been proven to survive all three Reidemeister moves, $V_K(t)$ is also a knot invariant.

We know that $V_K(t)$ is the Jones Polynomial because the corrected Kauffman Bracket can be transformed into the Jones Polynomial using a tool known as the Skein relation. The Skein relation functions similarly to the Kauffman bracket: It is a recursive formula that calculates the polynomial of a knot by relating it to simpler knots. The main difference between the two is that the Skein relation requires orientation, which is how it contains positive and negative crossings.

To build a Skein relation, we look at three knot diagrams that are identical everywhere except at one crossing:

L_+: The crossing is positive 

L_-: The crossing is negative 

L_0: The crossing is smoothed into parallel lines 

The Skein relation for the Jones polynomial is this:
$t^{-1}V(L_+) - tV(L_-) = (t^{1/2} - t^{-1/2})V(L_0)$

The Skein relation will break down every crossing (untangles the knot) by flipping or smoothing the crossing of a knot. Note that flipping the crossing (e.g. positive to negative) will always simplify/remove a crossing.

With the Kauffman bracket, we can build the Skein relation from scratch using the rules. If we apply Rule 2 to a positive crossing and a negative crossing, we get two equations:
$\langle L_+ \rangle = A\langle L_0 \rangle + A^{-1}\langle L_\infty \rangle$
$\langle L_- \rangle = A^{-1}\langle L_0 \rangle + A\langle L_\infty \rangle$

Note that L_∞ is the perpendicular smoothing, which only exists in the Kauffman Bracket rules. We have to eliminate this to get the Jones Polynomial.

Using basic algebra (multiplying the first equation by A, multiplying the second by A^{-1} and subtracting them), the L_∞ terms cancel out. This leaves us with the Skein relation for the Kauffman bracket:
$A\langle L_+ \rangle - A^{-1}\langle L_- \rangle = (A^2 - A^{-2})\langle L_0 \rangle$

However, this is just the regular bracket. To turn this into the Jones polynomial, we must apply our writhe correction and our substitution ($A = t^{-1/4}$).
Let w represent the writhe of the smoothed diagram L_0. Because L_+ contains one additional positive crossing compared to the baseline, its writhe is $w + 1$. Conversely, because L_- contains one additional negative crossing, its writhe is $w - 1$.
Substituting the writhe into our initial equation yields:
$A\Big[(-A^3)^{w+1}X(L_+)\Big] - A^{-1}\Big[(-A^3)^{w-1}X(L_-)\Big] = (A^2 - A^{-2})\Big[(-A^3)^{w}X(L_0)\Big]$

Notice that every term across the equation shares the factor of $(-A^3)^w$. By dividing the entire equation by this term, the writer is eliminated.
$A(-A^3)^1X(L_+) - A^{-1}(-A^3)^{-1}X(L_-) = (A^2 - A^{-2})X(L_0)$
Which simplifies into:
$-A^4X(L_+) + A^{-4}X(L_-) = (A^2 - A^{-2})X(L_0)$
To align with convention, we multiply the entire equation by -1 to invert the signs, yielding the corrected Kauffman bracket's Skein relation:
$A^4X(L_+) - A^{-4}X(L_-) = (A^{-2} - A^2)X(L_0)$
With the equation balanced and corrected, we can apply the substitution $A = t^{-1/4}$. This produces the corresponding transformations:

- $A^4 \to (t^{-1/4})^4 = t^{-1}$
- $A^{-4} \to (t^{-1/4})^{-4} = t^1 = t$
- $A^{-2} \to (t^{-1/4})^{-2} = t^{1/2}$
- $A^2 \to (t^{-1/4})^2 = t^{-1/2}$

Substituting these t variables into our equation, and changing the notation $V_K(t)$ to signify the Jones Polynomial, we arrive at the formula:

$t^{-1}V(L_+) - tV(L_-) = (t^{1/2} - t^{-1/2})V(L_0)$

This is the same as the Skein relation of the Jones Polynomial. By proving that the corrected Kauffman bracket with its substitution reduces into this equation, we prove that the invariant we constructed is the Jones polynomial.
Therefore, the Jones Polynomial is a knot invariant.

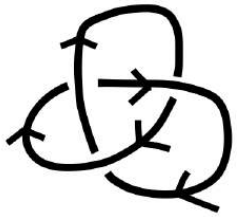
5.3 Unambiguity of the Jones Polynomial

Having established the state-sum formula for the Kauffman bracket, the proof of the Jones Polynomial being well-defined, or unambiguous, is rather straightforward.
The Jones polynomial is formulated by combining the Kauffman bracket with its writhe correction. Because the writhe of a given knot diagram is fixed, and the Kauffman bracket is unambiguous via the state-sum formula, their combined product is also well-defined. Therefore, establishing the Kauffman bracket via the state-sum formula intrinsically proves that the Jones polynomial is well-defined.

5.4 Example: The Jones Polynomial of the Trefoil Knot

We conclude by computing V(t) for the trefoil knot, illustrating Sections 3–5 on a single example.

Take the standard diagram of the trefoil with three crossings (Figure 3), oriented as in Figure 9, and suppose all three crossings are positive, so that w(D) = 3.



Step 1: The bracket. By the symmetry of the diagram, the 8 states fall into four classes according to how many crossings receive the A-smoothing and A^{-1} -smoothings:
Recalling the state-sum formula from proposition 3.2, we get these variables:

State classes:	**# states**	**α−β**	**loops |S|**
AAA	1	3	2

two A's, one A^{-1}	3	1	1
one A, two A^{-1}'s	3	−1	2
$A^{-1}A^{-1}A^{-1}$	1	−3	3

By the state-sum formula,

$$\langle K\rangle = A^3(-A^2-A^{-2})^1 + 3A^1(-A^2-A^{-2})^0 + 3A^{-1}(-A^2-A^{-2})^1 + A^{-3}(-A^2-A^{-2})^2$$
$$= (-A^5 - A) + 3A + (-3A - 3A^{-3}) + A^{-3}(A^4 + 2 + A^{-4})$$
$$= (-A^5 - A) + 3A + (-3A - 3A^{-3}) + (A + 2A^{-3} + A^{-7})$$
$$= -A^5 + (-A + 3A - 3A + A) + (-3A^{-3} + 2A^{-3}) + A^{-7}$$
$$= -A^5 - A^{-3} + A^{-7}$$

So, $\langle K\rangle = A^{-7} - A^{-3} - A^5$

Step 2: the normalized invariant. With w(D) = 3,

$$X(K) = (-A^3)^{-3}\langle K\rangle = -A^{-9}(A^{-7} - A^{-3} - A^5) = -A^{-16} + A^{-12} + A^{-4}$$

Step 3: the Jones polynomial. Substituting $A = t^{-1/4}$:

$$V(t) = -t^4 + t^3 + t$$

This matches the Jones polynomial of the right-handed trefoil.

6. Conclusion

The primary goal of this paper was to demonstrate how the Kauffman bracket, when corrected by the writhe, provides proof for both the invariance and unambiguity of the Jones polynomial.

The process began by recognizing the limitations of the Kauffman bracket. While it passed Reidemeister moves II and III, it failed Reidemeister Move I. Then, we introduced the writhe to act as a counterweight. This created the normalized bracket, $X(K)$, which is a knot invariant.

Once invariance was proven, the path to the Jones polynomial became rather straightforward. Through a substitution ($A = t^{-1/4}$), our invariant transformed into the Jones polynomial.

We proved the unambiguity of this invariant by using the state-sum formula in the Kauffman Bracket, which was transferred to the Jones Polynomial. In doing so, we proved that the Jones Polynomial is a well-defined knot invariant. Section 5.4 works through this whole process for the trefoil knot. The result, $V(t) = -t^4 + t^3 + t$ matches the known Jones polynomial and confirms that the writhe-correction formula is correct.

Works Cited

Adams, Colin C. The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots. W.H. Freeman, 1994.

Jones, Vaughan F. R. "A Polynomial Invariant for Knots via von Neumann Algebras." Bulletin of the American Mathematical Society, vol. 12, no. 1, 1985, pp. 103–11, https://doi.org/10.1090/S0273-0979-1985-15304-2.

Kauffman, Louis H. "State Models and the Jones Polynomial." Topology, vol. 26, no. 3, 1987, pp. 395–407, https://doi.org/10.1016/0040-9383(87)90009-7.

Roberts, Justin. Knot Knotes. Department of Mathematics, University of California, San Diego, 2015, math.ucsd.edu/~justin/Papers/knotes.pdf.

Orogeny of Mount Gyoja: Shaped through Opening Processes By Néo Petit

ABSTRACT

This research is part of a collective effort to clarify the extremely complex and controversial tectonic history of Japan. It aims to understand the orogeny of Mount Gyoja through an analysis of its rock outcrops; the study involves fieldwork, rock sampling and the detailed documentation of collision markers. My results confirm that the massif is composed primarily of granodiorite and include several findings: the coexistence of two types of granodiorites (standard and pink), the presence of a system of systematic joint families oriented towards the west and northwest, and a description of the massif's current general state. Based on these findings, I propose a scenario regarding the orogeny of Mount Gyoja.

KEYWORDS

Orogeny, Japan, Kameoka, Gyoja, Tanba, Tectonic, Bassin, Granodiorite, Lake, Erosion, Sea of Japan, Opening.

INTRODUCTION

Granitic rocks cover nearly 30% of Japan's land area and date primarily from the Mesozoic and Cenozoic eras (Nakajima et al., 2016). Many Japanese granites are part of the Upper Mesozoic granitic basement of the northwestern Pacific margin, which extends from Japan northeastward to Sikhote Alin and Chukotka, and southwestward to the Korean Peninsula and southern coastal China (Wu et al., 2022).

Proto-Japan formed nearly 700 million years ago when the supercontinent Rodinia fractured into several continental blocks—including Australia, East Antarctica, northern China, and Laurentia—which then gradually drifted apart (Isozaki & Yukio, 2023). Originally, Japan was part of the Eurasian continent—that is, it was attached to the South China block. However, 500 million years later, the Izagani Plate began subducting beneath the Eurasian Plate; as a result, the hydrated rock gradually brought the isotherm closer to conditions permitting partial melting of the Earth’s crust, between -200 and -100 Ma, promoting the gradual weakening of the back-arc volcanism of the North Chinese craton and initiating the process of separation. The primary magmas resulting from the mixing of these partial melts underwent various processes of fractional crystallization, giving rise to the Upper Cretaceous–Paleogene plutons of southwestern Japan (Wu et al., 2022). These magmas contributed to the growth of the continental crust by introducing new, non-recycled material. This process resulted in the separation (from -145 or -135 to -90 Ma), collision (from -90 Ma to -25 or -23 Ma) (Xu et al., 2024; Suga et al., 2020), and subsequent "final" separation of Japan from South China (from -25 or -23 to -14 Ma), leading to the opening of the Sea of Japan (Taverna, 1998). Specifically, between -25 and -23 Ma, a rift zone along the eastern coast of Eurasia created a ridge that initiated oceanic accretion, thereby isolating Japan, and opening the Sea of Japan. It thus began moving toward the oceanic

trench (Japan Trench) (Otofuji, 1996; Barnes, 2003), which marks the subduction zone between the Eurasian, Pacific, and Philippine Sea plates (Gianni, 2025).

The Sanyo Belt (山陽) formed during the Late Cretaceous and is dated to between -100 and -70 Ma (Li & Ishihara, 2013). It lies on the northern side of the Median Tectonic Line, in southwestern Japan (Ishihara et al., 1996), and contains igneous rocks dating from the Late Cretaceous–Paleogene. It comprises several groups, including the Tanba Group (丹波), which encompasses the Kameoka Basin (亀岡), where Mount Gyoja (行者) is located.

The Tanba Zone includes the vast mountainous region in the center of the prefecture. Its southern boundary is marked by the Hozukyo Gorge (保津峡), and its northern boundary by the edge of the Kyoto Basin (京都), which extends north of Mount Hiei (比叡). The Tanba Group consists of a succession of oceanic plate strata rich in adakite-type volcanic rocks formed by the melting of subducted oceanic crust (Li & Ishihara, 2013).

The Kameoka Basin is a narrow, elongated lake basin, approximately 7.5 miles long, oriented northwest-southeast, and about 1.8 miles wide, oriented northeast-southwest. It originated from tectonic subsidence approximately 4 million years ago, at the beginning of the Pliocene epoch (Kawabe, 1989). It was formed by block movements and subsidence caused by the tilting of faulted blocks a few miles in an east-northeast direction. A reservoir, located at about 919 feet above sea level, subsequently occupied the basin until the Pleistocene, and analysis of the topography and geology suggests that the area was also covered by a lake or marsh on several subsequent occasions. According to Okada et al. (2005), the area contains a fault system with a total length of approximately 8 miles extending northwestward. The segment closest to Mount Gyoja (“F1”) is a steeply dipping eastward reverse fault. It is approximately 1.9 miles long and located less than 2 miles from the mountain. It has played a major role in the formation of the basin's underground structure since approximately 870 Ka and is considered active.

Mount Gyoja was formed by a granitic intrusion approximately 93 million years ago, during the Late Cretaceous (Morishita, 1991). It contrasts with the other geological areas of the Tanba Group because it is the only granitic massif in this region—and, moreover, an isolated one. According to Hayase & Tsutsumi (1956), it contains only a few quartz dikes. Mount Gyoja is in contact with Paleozoic strata to the north and Mesozoic strata to the west and is overlain by Cenozoic strata to the south and east.

OBJECTIVE

The objective is to determine the orogeny of Mount Gyoja based on its outcrops, their composition (via visual inspection), and any possible collision markers.

MATERIAL AND METHODS

Strategy and Site Selection

Mount Gyoja is a relatively small, dome-shaped massif (Fig. 1), but a study based on a meticulous and systematic grid would be time-consuming and potentially dangerous on steep slopes. To increase my chances of finding rock outcrops and to make my time investment worthwhile, I used the map (Fig. 2) illustrating the article by Hayase & Tsutsumi (1956, Fig. 1), who had, among other things, tried to locate the limits of the massif, because it indicates where the outcrops that allowed their sampling were located. However, as this map was very rudimentary, I had to determine approximate GPS coordinates using the Google Map and Geological Survey of Japan websites and Garmin BaseCamp.

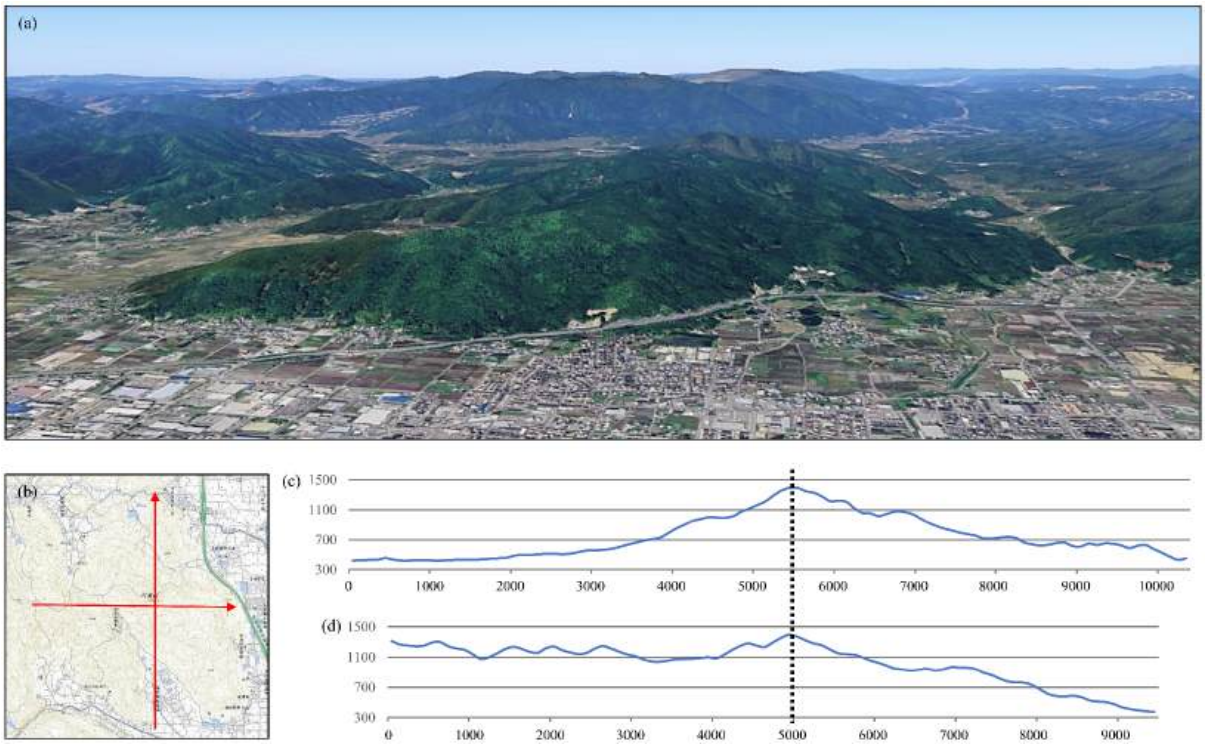


Figure 1: Appearance of Mount Gyoja. (a) Aerial view of the mountain from the east; (b) Map showing the area occupied by the mountain; the two axes (south-north and west-east) passing through the summit of Mount Gyoja are used for the topographic representations. Topography of the mountain along two axes: (c) S-N (10560 feet long) and (d) W-E (9134 feet long), generated using data provided by Garmin BaseCamp. The x- and y-coordinates are in feet. Both topographic representations are drawn to the same scale, with a 1:1 ratio between the x- and y-coordinates. The vertical dotted line indicates the position of the summit.

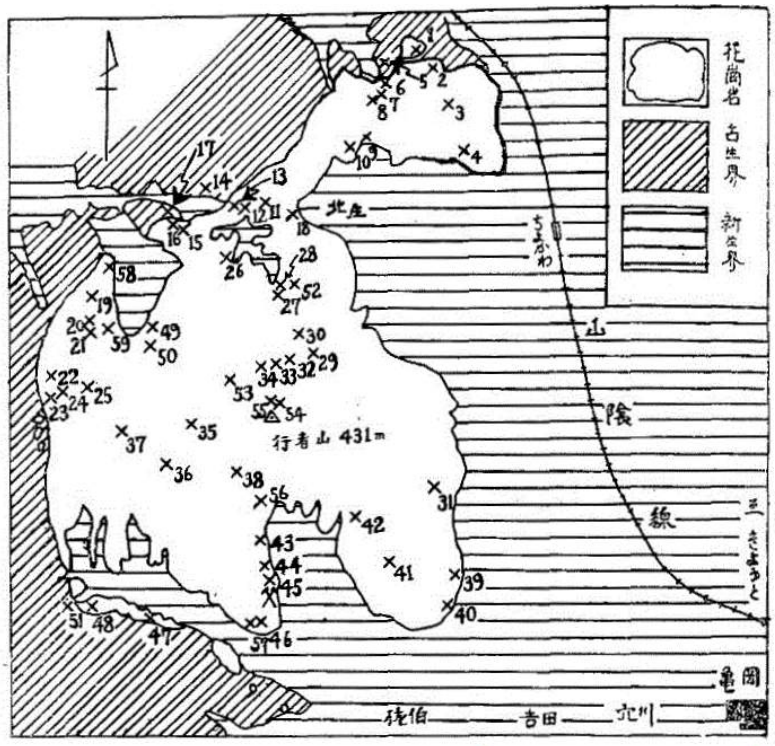


Figure 2: A rough map from Hayase & Tsutsumi's article showing their rock sampling sites.

I then planned my routes and attempted to locate these sites. This approach proved highly successful, though imperfect. Since 1956, some areas have become inaccessible, and many outcrops have likely vanished due to urbanization or are now obscured by vegetation and sediment. Conversely, I added additional sites during my explorations to document the widest possible area.

Equipment: GPS navigator Garmin Etex 30 with Japan Topo 10M Plus V3 map; software: Garmin BaseCamp (Version 4.4.6); websites: Geological Survey of Japan (https://gbank.gsj.jp/geonavi/), Google Map (https ://www.google.com/maps/); the map from Hayase & Tsutsumi (1956).

Description of the Area

The area I studied consists of three zones (Fig. 3). First, there is a zone located north of Mount Gyoja, near Yagi Station (Sagano Line, Japan Railways), separated from the mountain by National Route 452, with an area of approximately 0.3 square-miles (maximum dimensions: 0.75 by 0.6 miles) and a maximum elevation of 690 feet. I named it "Northern Outer Zone" for this research. Second, there is Mount Gyoja itself; it covers an area of approximately 2.6 square-miles (maximum dimensions: 1.7 by 2 miles), has a maximum elevation of 1,476 feet, and is dome-shaped (Fig. 1). Finally, there is a zone located south of Mount Gyoja, near Namikawa Station (Sagano Line, Japan Railways); it is separated from the mountain by National Route 372, covers an area of 0.3 square-miles (maximum dimensions: 0.75 by 0.1 miles), and has a maximum elevation of 525 feet. I named it "Southern Outer Zone".

These three areas are identified by Geological Survey of Japan as composed of granodiorite. They are contiguous zones and the only granitic areas of this type within the Tanba Group. They likely belong to the same intrusion and were considered by Hayase & Tsutsumi (1956) to be part of Mount Gyoja.

Rock Sampling and Site Documentation

For each sample, I followed this procedure: (1) Took several photos (using my iPhone and camera) of the site, including one with a reference target, (2) Recorded the GPS coordinates and elevation using the navigation app, (3) Collected 2 or 3 samples using a geologist's hammer, (4) Numbered and bagged each sample, and (5) Took written notes or recorded voice memos on the iPhone.

Equipment: small scale bar (white graduated 6-by-3-inch square) for photographs, geologist's hammer, safety glasses, chalk for numbering samples, notebooks, individual storage bags for each sample, Iphone i15, Olympus TG-6 camera, GPS navigator for geolocation.

Documentation of Collision Markers

For each marker, I followed this procedure: (1) Took several photos (using my iPhone and camera) of the site, including one with a reference target, (2) Recorded the GPS coordinates and elevation using the navigation app, (3) Collected 2 or 3 samples using a geologist's hammer, (4) Numbered and bagged each sample, (5) Took written notes or recorded voice memos on the iPhone, and (6) Made a field sketch.

Additional equipment: Bosch GLM50 rangefinder, compass, FieldMove Clino (version 3.1.3) installed on the iPhone.

Rock Sample Processing

The resources available to me for sample analysis were very limited. Lacking the means for advanced testing, I opted to use macrophotography of freshly polished flat surfaces to identify the samples' structures and to perform hydrochloric acid tests to identify limestone. I also photographed the slices after hydrating them to better reveal components obscured by quartz. Data from the article by Hayase & Tsutsumi (1956) regarding the composition of samples collected in the 1950s will serve as a reference. I assume that the granitoid rocks I collected—either in whole or in part—share the same characteristics as theirs, thereby obviating the need for such advanced testing.

Rock Cutting and Polishing

The cutting and polishing work was carried out by my father due to the dangers involved in handling the grinder and the respiratory risks associated with inhaling the fine dust generated during polishing. A portion of each sample was also set aside for the hydrochloric acid test.

Equipment: Makita 95533BL angle grinder, Kai Kitchen AP0305 whetstone.

Macrophotography

The photos of the samples were taken with the TG-6. The samples were first weighed and photographed in their original condition from several angles. After cutting and polishing, I took macrophotographs of the samples. For the granite samples, I also took photographs of the hydrated sections. In fact, in the hope of highlighting components of the granite rocks hidden by the quartz, I applied water to the sections of all my samples.

Analysis of the Modal Composition of the Samples

To determine the proportions of the components in granitic rocks, I used a visual approach based on proportional composition charts from Terry & Chilingar (1955) and charts produced by John A. Stevenson (2012) on his website (https://all-geo.org).

Acid Test

A small piece of each sample was exposed to several drops of 10% hydrochloric acid. The presence or absence of effervescence and/or whitening of the sample within 5 seconds was noted. For samples suspected of being sedimentary, a larger amount of acid was applied to the entire sample, and I observed the sample for 10 seconds.

Equipment: bottle of 10% hydrochloric acid (Matsushima Chemical), pipette, protective gear.

RESULTS

During the month of May 2026, I spent 4 days exploring and covered about 45 miles of mountain hiking and mountaineering. I likely reached 27 of Hayase & Tsutsumi's 59 sites (1956) and documented 16 new sites. Forty-seven viable samples were collected, totaling approximately 18 pounds, and 42 sites were documented (Fig. 3; Extended data, Table 1), including 5 collision markers.

Acid Test Results

All samples were tested with 10% hydrochloric acid (HCl). No visible reactions occurred (neither effervescence nor whitening). Therefore, none of the samples contain calcium carbonate ($CaCO_3$) (Extended data, Table 1).

Rock Identification Results

The samples are divided into three categories: metamorphic rock, basalt, and granodiorite (Extended data, Table 1).

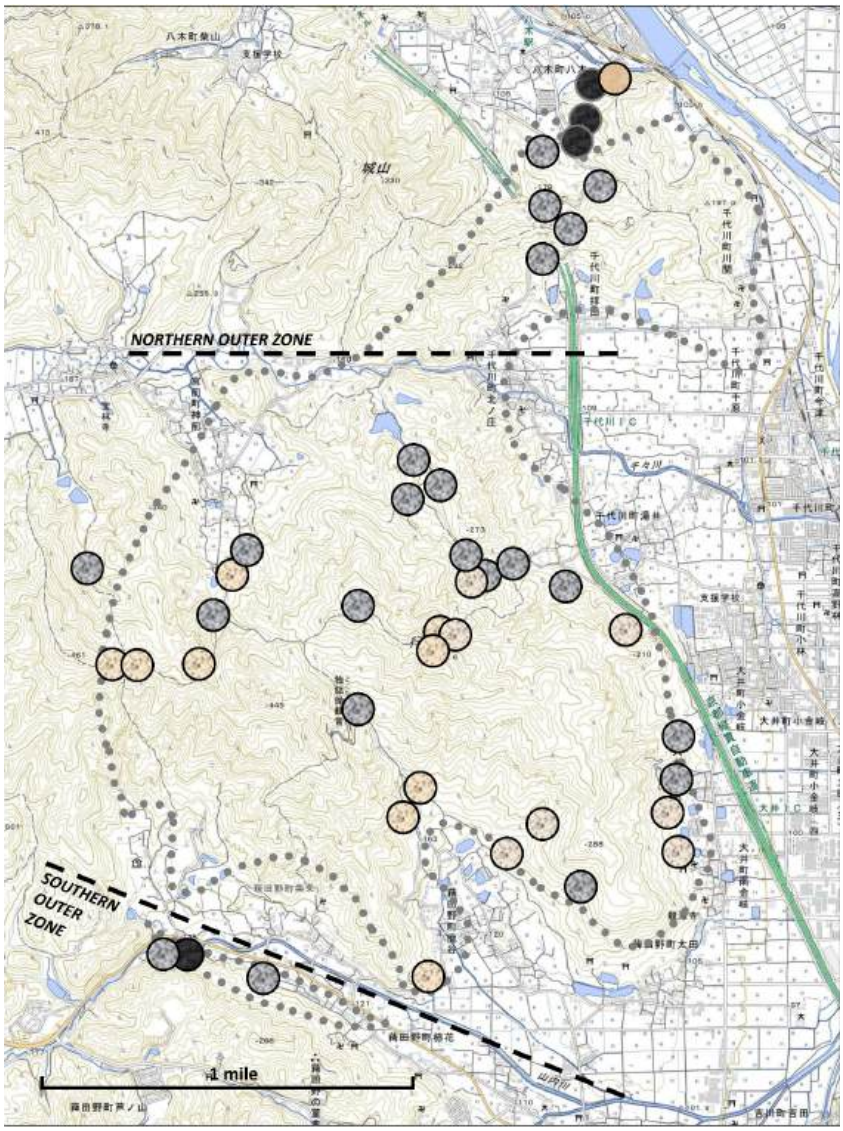


Figure 3: Locations of sampling sites and types of samples collected. Basalt samples are shown in black; metamorphic samples in beige; granodiorite samples without pink feldspar in gray; and granodiorite samples with pink feldspar in pink. The granitic zone, based on Geological Survey of Japan's site is indicated by dotted lines. Long black dashed lines mark the boundaries between the official Mount Gyoja and the outer zones.

Metamorphic Rocks

Three samples were collected (samples 3a, 3b, and 3c, Extended Table 1) from collision marker A, located in the Northern Outer Zone.

My photos and the information available on the geological maps suggested that my samples consisted of sedimentary rocks. The appearance of these samples suggests pelite with chert inclusions (Fig. 4a and 4b). However, the lack of a reaction to the acid test ruled out this hypothesis. It would therefore appear to be a fine-grained metamorphosed sedimentary rock, likely pelitic schist, a metamorphosed form of pelite. Nevertheless, only advanced laboratory

tests would allow for a definitive identification. For the purposes of this analysis, I will therefore refer to it as a "metamorphic rock," without being more specific.

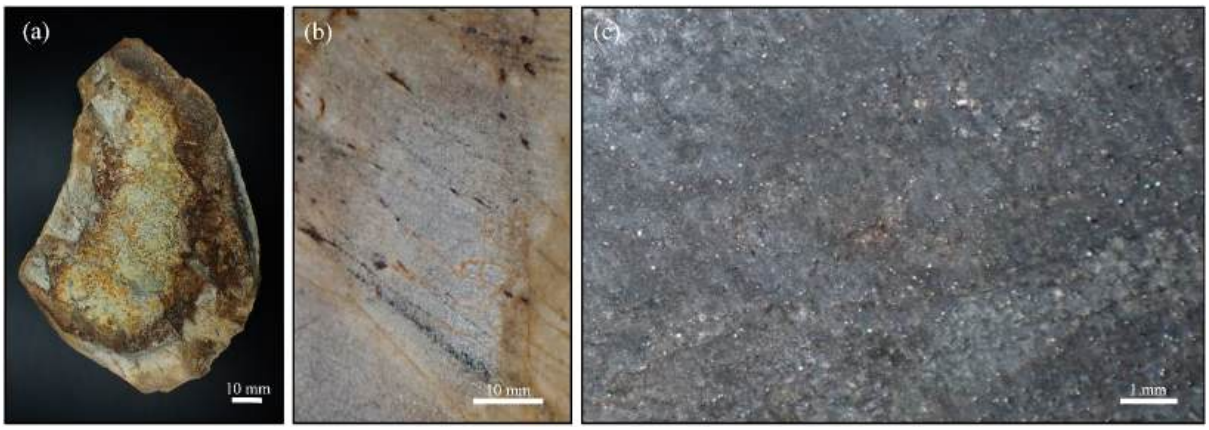


Figure 4: Metamorphic rock (a) original sample and (b) polished slab, and (c) basalt slab.

Basalt

All basalt samples (samples 2a, 2b, 4, 5, and 41, Extended data, Table 1) were collected from the outer zones. To the north, three sites are distributed along a straight line approximately 600 feet long, and two of the four samples come from marker B.

Five samples consist of a black to dark gray, aphanitic rock containing many small white or light gray microlites distributed relatively uniformly, evoking "a starry night" (these may be plagioclase feldspar) (Fig. 4c). On the other hand, two samples contain goethite at the periphery, and exhibit clusters of goethite, accounting for 20 to 30% of the section. These features point toward an iron-rich tholeiitic basalt. However, here again, only advanced laboratory tests—notably petrographic microscopy and chemical analysis to determine the proportions of plagioclase feldspar and pyroxene, as well as the calcium, silica, and sodium contents—would guarantee identification of the specific type of basalt and its origin.

Granodiorite

All other samples (N=39) are granitoids (Fig. 5a). They exhibit a phaneritic texture containing quartz, plagioclase feldspar, alkali feldspar, and mica (biotite), like Hayase & Tsutsumi's samples (1956). Some of them are weathered: rust-colored halos surround some of the biotite crystals (N=6) and/or the sample exhibits an ochre-brown color on its surface; several of the plagioclases are weathered and appear to have been replaced by a clay mineral (N=4). The remaining samples (N=30, or 63.8%) are relatively fresh. A few samples are porphyritic (N=4). The others (N=35) are equigranular and lack stratification, banding, or porosity. The texture consists of interlocking crystals of translucent to off-white quartz and a significant proportion of plagioclase feldspar (Fig. 5). The quartz generally appears translucent, light gray, or off-white in the granodiorite, and one has the impression of being able to see through the surface when the section is soaked in water. Furthermore, a visual analysis of the modal composition of the

undegraded, equigranular samples based on proportional composition charts indicates that the proportion of dark mafic minerals is less than 15%. In addition, the proportion of plagioclases is significantly higher than that of alkali feldspar in most of the samples. For all these reasons, I can confirm that all these samples correspond to granodiorite, which corroborates the data from the geological maps.

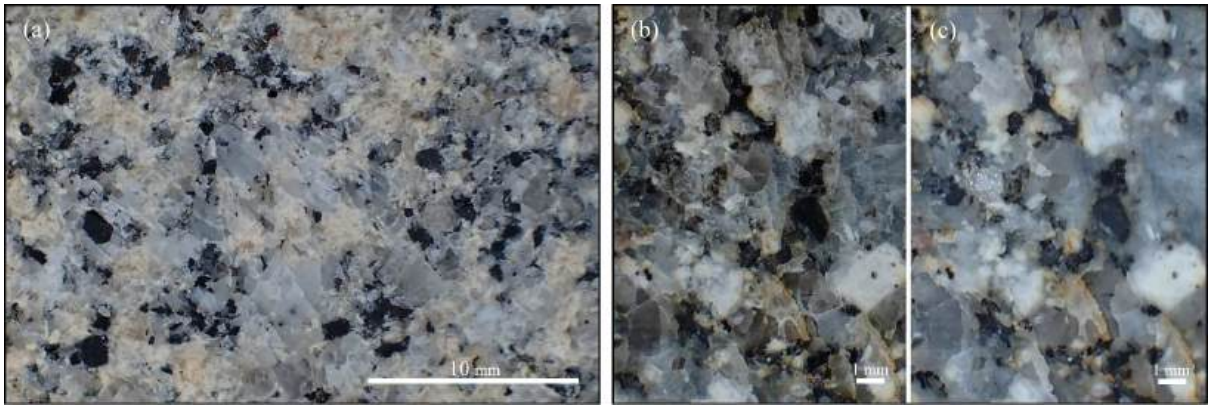


Figure 5: Large slab of standard granodiorite, shown in a wide view (a) and in close-up, with the same slab (b) dry and (c) hydrated.

Nearly half of the samples (48.7%) contain salmon-pink, dark pink, or red alkali feldspars (Fig. 6). In three cases, I observe only a few clusters; in the others, they occupy extensive areas. When I hydrate the cross-section of several samples, the crystals even take on a crimson hue (Fig. 6b, 6c), particularly those from the summit. I also note that the samples from the Northern and Southern Outer Zones do not contain pink granodiorite.

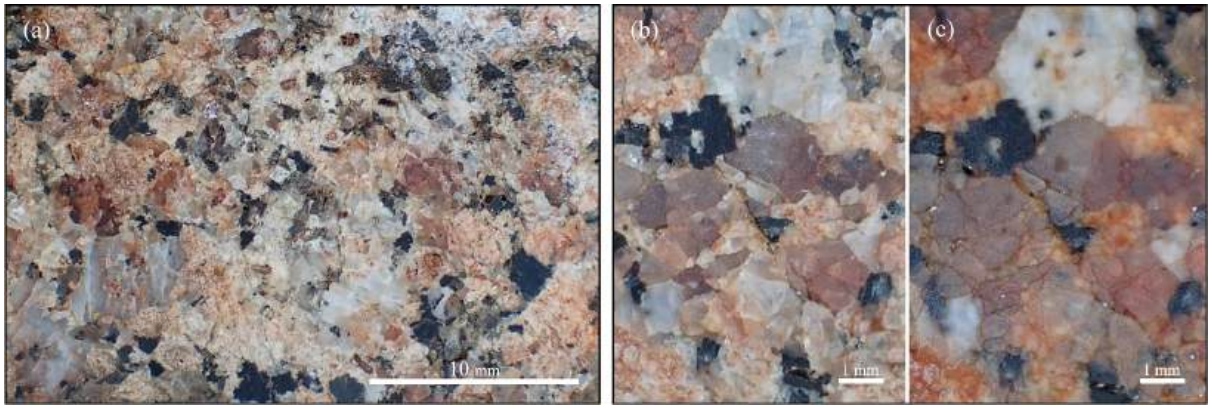


Figure 6: Large slab of pink granodiorite, shown in a wide view (a) and in close-up, with the same slab (b) dry and (c) hydrated.

Description of Collision Markers

I have identified five collision markers: three in the northern outer zone (A, B, and C) and two on Mount Gyoja itself (D and E) (Fig. 7).

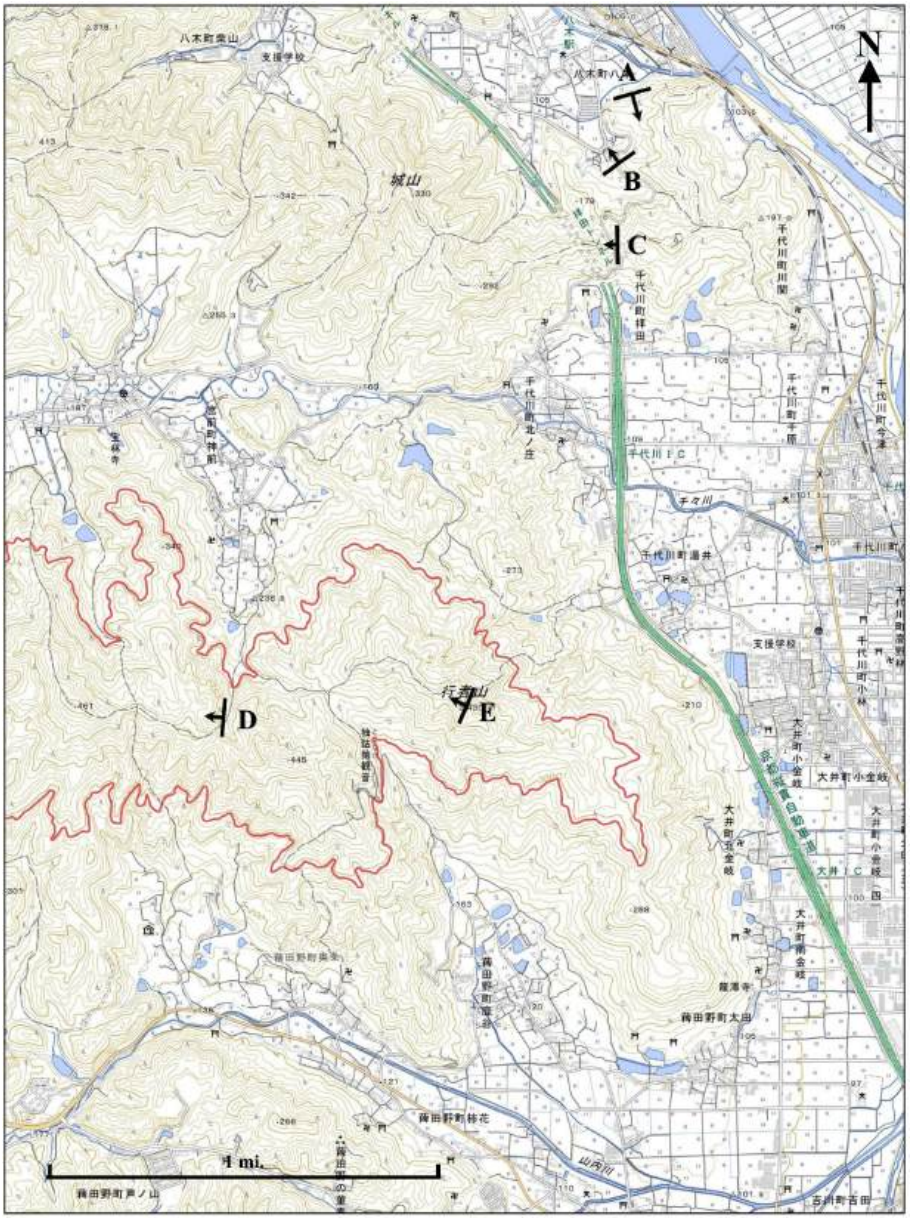


Figure 7: Location of collision markers with dip, and a red line marking the 919 feet elevation.

<u>Collision Marker A</u> (Fig. 8)

It is composed of metamorphic rock (samples 3a, 3b, 3c, Extended data, Table 1). It is located on a hillside at an elevation of 440 feet, in the heart of a cypress forest, and consists of a series of 9 outcrops spread over a length of approximately 65 feet, facing a plain. The segments are separated from one another by 3 to 30 feet of plant debris. Three different samples showed that they consist of metamorphic rock, likely a pelitic schist.

Each segment is angular, often prism-shaped, and consists of parallel strata of comparable thickness. The segments measure approximately 3 to 7 feet long and 0.1 to 3 feet wide.

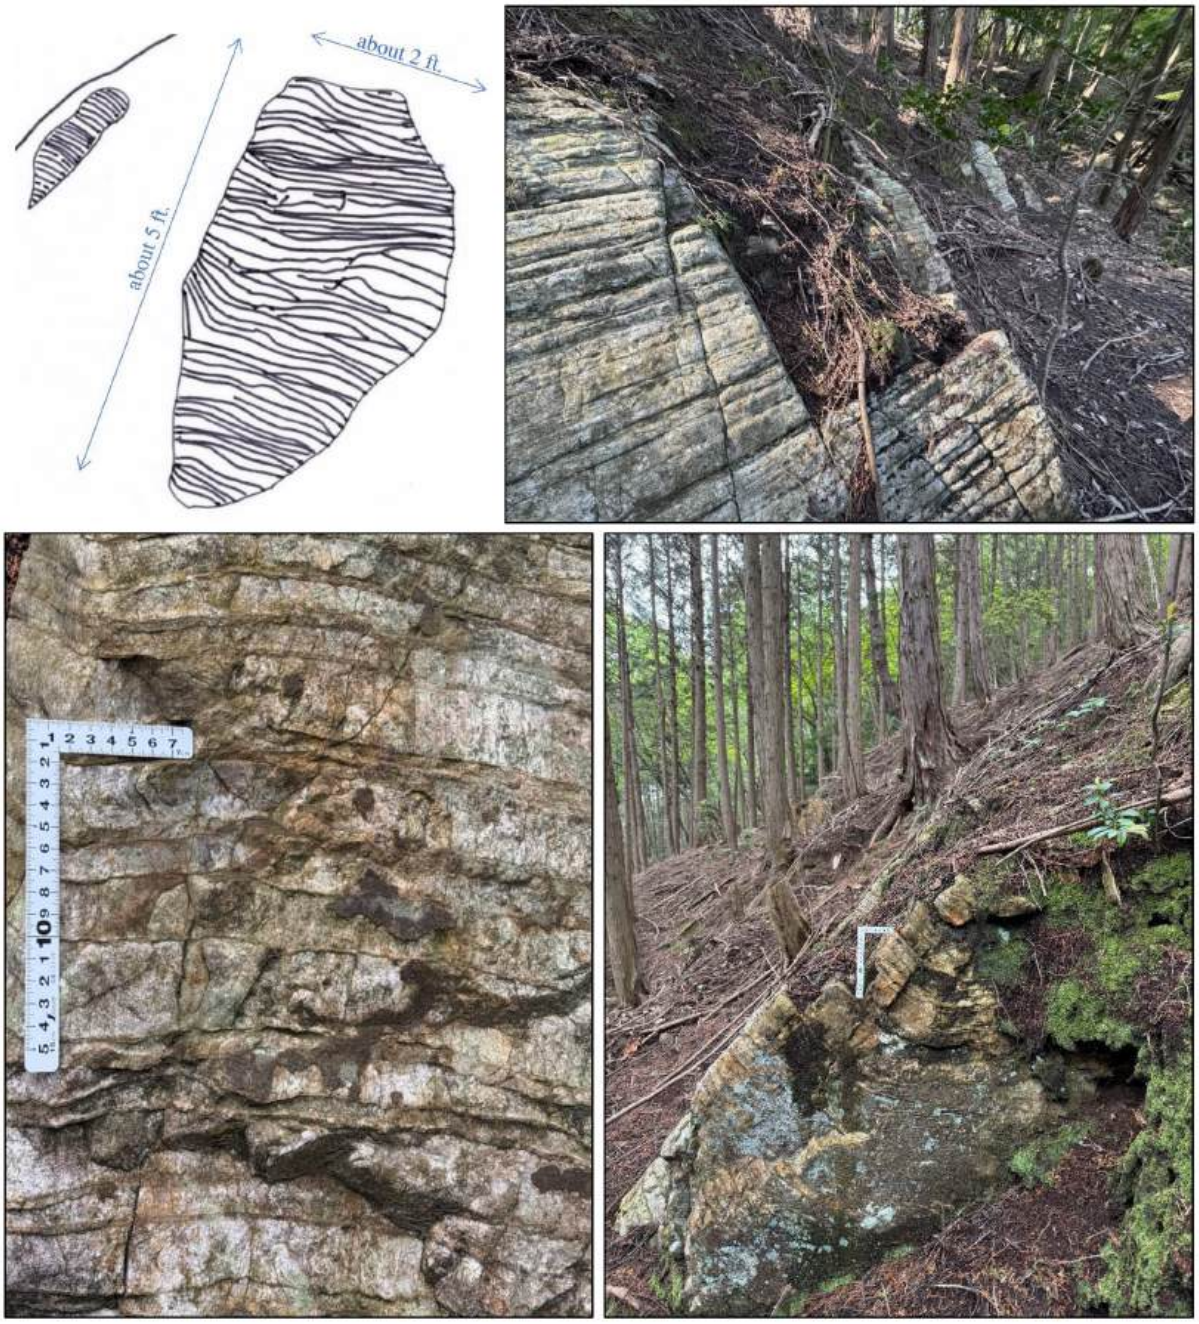


Figure 8: Views of Marker A: (a) field sketch, (b) photo of 7 of the 9 segments, (c) close-up photo of the strata, and (d) photo of the overturned synclinal fold.

On average, the dip is oriented towards the south (N=5; $\bar{x}$=164,17; δ=13,15), with an inclination of 27,33° (δ=7,84) (Extended data, Table 2).

The easternmost outcrop reveals that this marker is an overturned synclinal fold (dipping at 45°). It measures approximately 28 inches in height, 28 inches at the base, and 40 inches along the fold limb. Consequently, dip measurements on the outcrops were taken on the limbs preceding this hinge; the fold axis trends northward, opposite to the direction of the dip.

<u>Collision Marker B</u> (Fig. 9)

It is located at an elevation of 406 feet and is composed of basalt (samples 2a, 2b, Extended data, Table 1). Measuring 13 feet high and 72 feet wide, it is situated along the roadside. Sampling confirms the material is basalt. It consists of a set of systematic joints forming at least 10 blocks, each approximately 0.7 to 1.5 feet thick. The rocks are heavily weathered, coated with clay, and largely obscured by vegetation and sediment. The lower and upper boundaries are diffuse.



Figure 9: Views of Marker B: (a) field sketch, (b) photo of one section of the marker, and (c) photo of another section of the marker.

On average, the dip is northwest (N=3; $\bar{x}$=324; δ=7), with an inclination of 53,33° (δ=1,53) (Extended data, Table 2).

<u>Collision Marker C</u> (Fig. 10)

It consists of standard granodiorite (samples 8a and 8b, Extended data, Table 1). The marker is located along the roadside at an elevation of 555 feet in a bamboo grove; it is 59 feet wide and approximately 78 feet high. A core sample confirmed that it consists of granodiorite. This is a family of systematic joints consisting of five compartments ranging from 27 to 40 inches thick. The lower and upper boundaries are diffuse. The rocks are heavily eroded and surrounded by granitic sand, and the breccias are intensely clay-rich. The early stages of granitic chaos formation are evident.

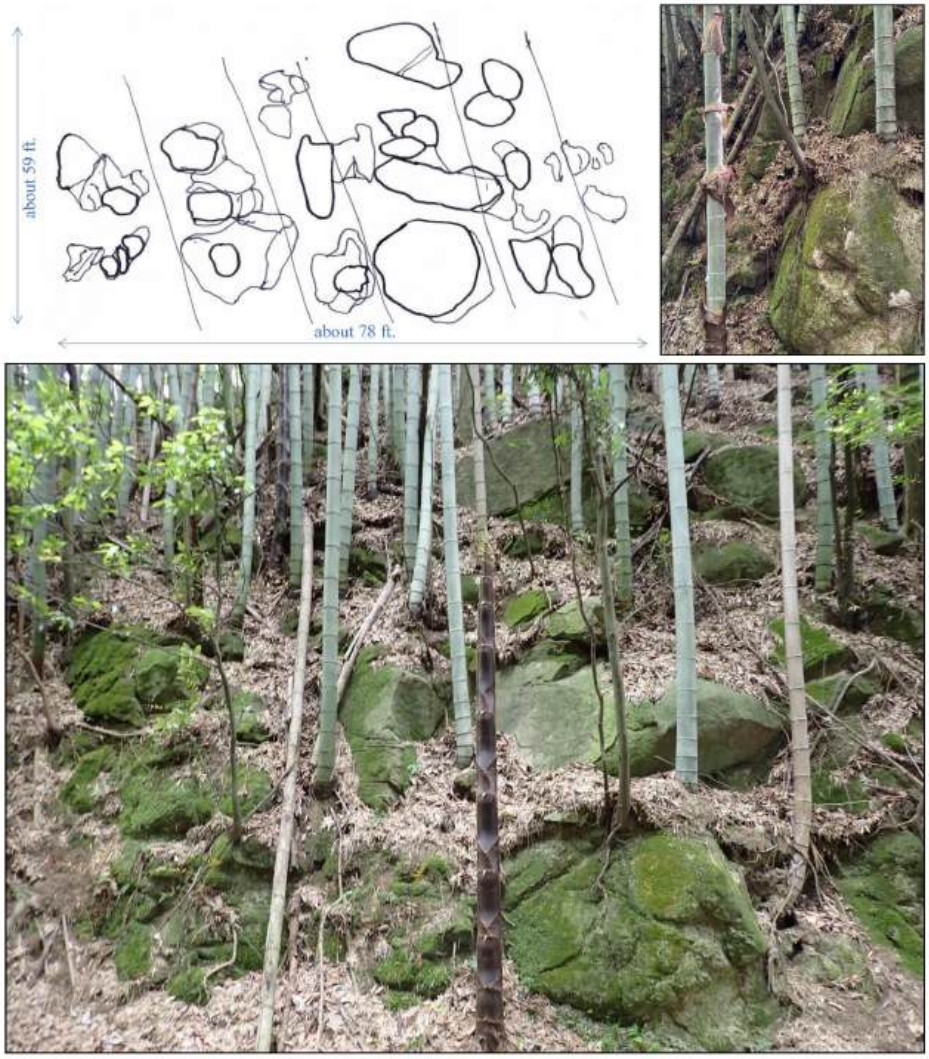


Figure 10: Views of marker C: (a) field sketch, (b) photograph of a section of the marker, and (c) joints viewed from a different angle.

On average, the dip is westward (N=5; $\bar{x}$=269; δ=13.05) with an inclination of 76° (δ=3) (Extended data, Table 2).

<u>Collision Marker D</u> (Fig. 11)

It is composed of pink granodiorite (sample 33, Extended data, Table 1). It is located 0.7 miles west of the mountain summit at an elevation of 1,148 feet, near a cypress forest and at the edge of a clearing; it measures approximately 22 feet in length and 9 feet in width, with an exposed height of about 9 feet. Sampling confirmed it to be granodiorite.

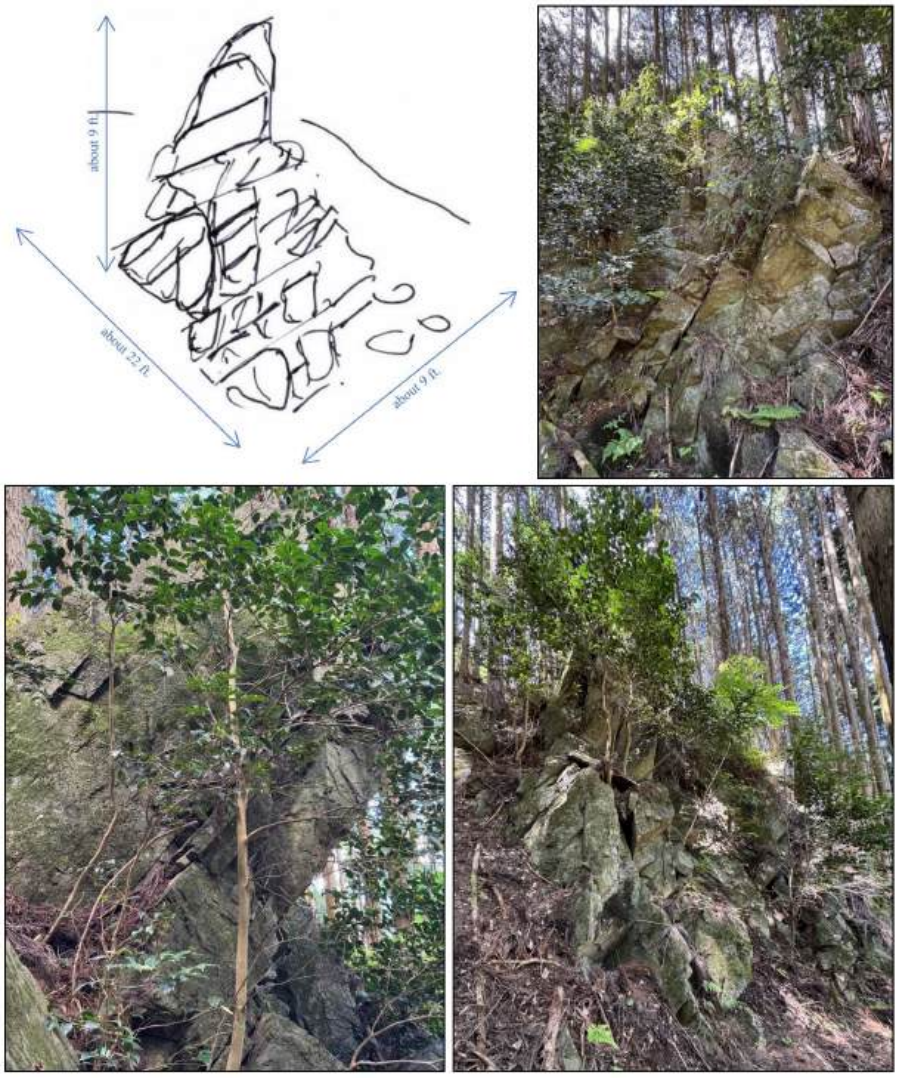


Figure 11: Views of Marker D: (a) field sketch, (b) close-up photo of joints, (c) overall view from below, and (d) photo of the rear section.

The formation resembles a whale breaching the water. The rock is fresh, with sharp, straight edges. It is a family of systematic joints comprising more than 20 blocks, ranging from 4 to 20 inches in thickness. However, several trees have grown within the separated sections.

On average, the dip is westward (N=2; $\bar{x}$=275; δ=1,42), and the inclination is 59,50° (δ=0,70) (Extended data, Table 2).

<u>Collision Marker E</u> (Fig. 12)

This is the most impressive landmark. It is composed of pink granodiorite (samples 36a and 36b, Extended data, Table 1). It is located at an elevation of 1,342 feet, near the summit of Mount Gyoja. Incidentally, it is part of a Shugendo temple that includes altars and other buildings.



Figure 12: Views of marker E: (a) field sketch, (b) close-up photo of a joint, (c) photo from above, and (d) overview photo from below.

This family is part of a system of distinct, systematic joint families extending over several hundred feet. This system consists of three parts (upper, middle, and lower), and marker E corresponds to the middle part. It is the largest and, in my view, representative of this system; the others have the same appearance and, overall, an apparently similar dip. Marker E is approximately 100 feet long, 40 feet high, and 49 feet wide. It is a family of systematic joints composed of 18 compartments of granodiorite ranging from 8 to 8.3 feet in thickness. The rock is fresh, the joints are straight, and there is no displacement or filling. However, several trees have grown within the jointed sections.

On average, the dip is northwest (N=4; $\bar{x}$=293,25; δ=14,8), and the inclination is of 43,50° (δ=6,35) (Extended data, Table 2).

Synthesis

I discovered a system of joint families. Three families are composed of granodiorite (C, D, and E) and one of basalt (B). They are all oriented towards the west or northwest with moderate to steep angles, ranging from 38° to 79°.

State of Deterioration of the Outcrops

The topographic profile of Mount Gyoja resembles that of an ancient, eroded massif. The massif is forested and has relatively few outcrops. They appear to be more numerous in the north and east. Most of the outcrops I observed consist of rounded boulders in areas rich in granite sand. Four out of 42 outcrops are even limited to slightly domed surfaces less than 12 inches high. Furthermore, many outcrops show signs of weathering (fractures, flaking, decay, etc.). However, there are also large outcrops with sharp, straight edges showing no signs of erosion, but I note that all those I found are located at an elevation above 919 feet (Fig. 7), —that is, above the level of the lake that once occupied the Kameoka Basin (Kameoka City History Compilation Committee, 1995). Specifically, the collision markers located below this line (A, B, and C) are all heavily eroded. In contrast, those located above it (D and E) are not (Fig. 7), and I note that they are the only markers composed of pink granodiorite.

ANALYSIS

My samples confirm that the massif is composed of granodiorite. It is likely that the granodiorite intrusion was formed during an upwelling of partially melted peridotite during the penultimate phase of the Yanshan orogeny, a phenomenon described by Suga & Yeh (2022), and to my knowledge, this is the only available explanation. My samples also reveal that nearly half of the outcrops on the mountain itself contain pink alkali feldspar, but lacking suitable tests, I am unable to explain why, or to determine whether this could have influenced the rock's resistance to water erosion.

I identified five collision markers. The first, located at the edge of the northern outer zone, consists of metamorphic rocks. This fold, which is now heavily eroded, may have been caused by contact metamorphism during the ascent and cooling of the intrusion; this type of curved structure is, in fact, due to the ductile deformation of sedimentary layers under compressive stress. The other markers are families of systematic joints. One consists of basalt despite its location in the northern outer zone. It is likely that this marker and the other basalt outcrops I have documented are remnants of the accretionary strata within which the intrusion occurred. The other three families consist of granodiorite; none contain veins, and two consist of fresh rock.

This type of joint is typical of stresses that stretch a rock layer to the point of exceeding its tensile strength, eventually causing it to fracture. In this case, the rock fractures in a plane parallel to the maximum principal stress and perpendicular to the minimum principal stress (the direction of stretching). Low tectonic stresses are required, and joints form readily. The appearance, dimensions, and arrangement of joints can vary considerably, and they are found in

very different tectonic settings, so the precise origin of the stresses that created them is often uncertain and sometimes controversial. In this case, given the tectonic history of this region since the intrusion, I considered three possible origins. First, there is the eastward opening of the Sea of Japan, according to the "double-door" model (Otofuji, 1996; Barnes, 2003) (Fig. 13).

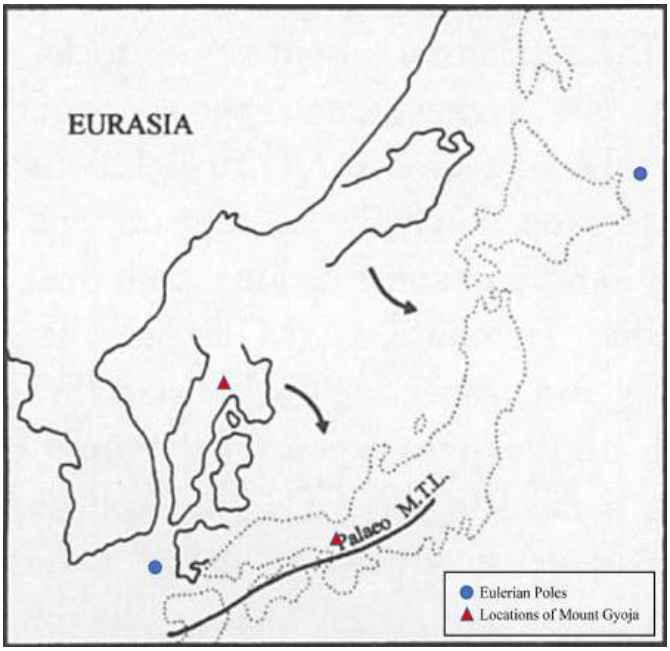


Figure 13: Position of South Japan and North Japan at the continental margins before the opening of the Sea of Japan basin and the outward rotation of the landmasses (after Otofuji, 1996, Fig. 5a; after Barnes, 2003, Fig. 8b).

The direction of rotation around the Eulerian pole located to the southwest is consistent with the direction of the joints and the expansion process involving tensile fracturing. However, I am surprised by the freshness of outcrops D and E and the absence of veins in all joint families if they do indeed date to the Lower Miocene. Second, there is the tectonic subsidence responsible for the opening of the basin, approximately 4 million years ago (Kawabe, 1989), in an east-northeast direction (Fig. 14).

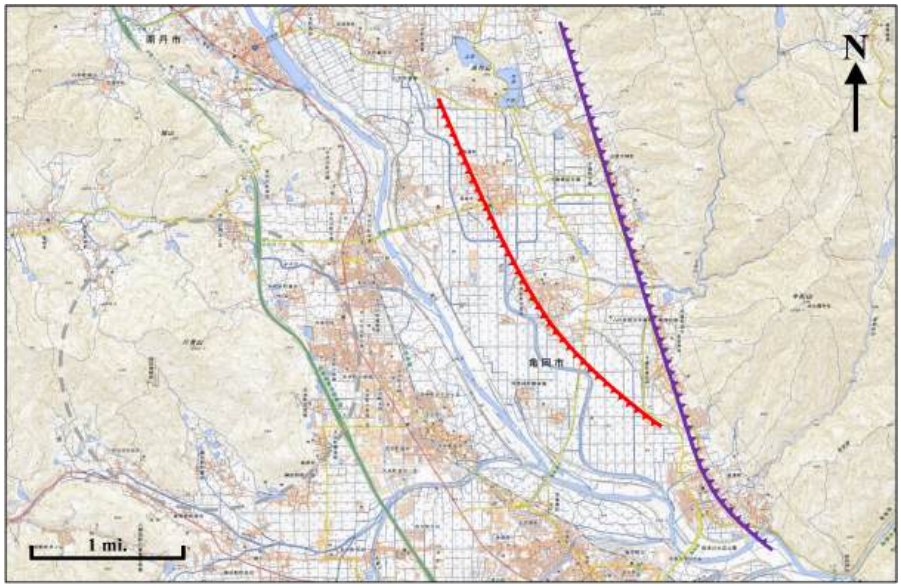


Figure 14: A portion of the Kameoka Basin located near Mount Gyoja (circled in gray dotted lines). The "F1" fault (after Okada et al., 2005, Fig. 2, 3) is shown in red, and the subsidence fault is shown in purple (Kawabe, 1989, Fig. 1, 3a).

The margins of the collapse pit during the subsidence, where the intrusion was necessarily located, were subjected to tension parallel to the direction of subsidence; the joints should therefore be perpendicular to this direction. However, although the stretching phenomenon is consistent with the type of joints I found, their directions are not. Thirdly, there is the "F1" fault (Okada et al., 2005) located in the immediate vicinity of the mountain (Fig. 14). However, being a reverse fault with a steep dip to the east, the tectonic pressures on the mountain located west of the fault are necessarily compressive and therefore cannot be the origin of the joints. Consequently, based on the current state of knowledge regarding the tectonic past of this region, it appears that the joints I found resulted from stresses related to the opening of the Sea of Japan 25 to 23 million years ago.

The current dome shape of Mount Gyoja suggests that it was originally a laccolith. Although relatively young, the massif exhibits all the characteristics of an ancient, highly eroded massif (Hayase & Tsutsumi, 1956). The lake, which occupied the basin and in which the intrusion was submerged for one to two million years (Kameoka City History Compilation Committee, 1995), is necessarily partly responsible for its accelerated erosion.

Joints facilitate the weathering of granite, which accelerates their disappearance when submerged. Consequently, it is unlikely that any jointed outcrops located below the 919 feet boundary have survived to this day, and it is likely that many others existed and contributed to the accelerated erosion of the massif.

Finally, I have no way of knowing the volume eroded away; however, the northern and southern outer areas are potential indicators of the area originally occupied by the laccolith.

CONCLUSION

To explain the orogeny of Mount Gyoja, I propose the following model. Approximately 93 million years ago, an intrusion composed of magma from the penultimate phase of the Yanshan orogeny formed a granodiorite laccolith within sedimentary strata and accretionary complexes. Nearly 70 million years later, the opening of the Sea of Japan triggered the formation of systematic joint families at various elevations. Then, about 4 million years ago, tectonic subsidence created the Kameoka Basin, on the edge of which the massif is located; a lake subsequently formed there and remained until the Pleistocene, causing accelerated erosion of the mountain. Finally, about 2 million years ago, the lake drained, the erosion slowed down and became primarily driven by precipitations.

DISCUSSION

To continue this research, I would like to conduct a systematic study using a fine grid pattern to identify any additional collision markers on this massif, as well as on massifs located in the same region. To this end, using a swarm of drones (Peacock & Corke, 2020; Vansutre et al., 2025) equipped with LiDAR (Lo et al., 2021) and guided by artificial intelligence would be ideal for locating them, minimizing time spent and reducing risks. The use of LiDAR would also allow for precise measurement of the dimensions of the outcrops and their dips under challenging conditions.

Cutting and polishing the samples allowed for a relatively reliable identification of plutonic rocks, and hydrochloric acid tests ruled out the possibility of sedimentary rocks. To complete this research, I would like to conduct advanced tests on my samples (such as energy-dispersive X-ray spectroscopy and scanning electron microscopy) to confirm my results and study the physicochemical characteristics of the pink granodiorite.

ACKNOWLEDGEMENTS

I would like to thank Ms. Poppeschi, geologist, and science teacher at the Lycée Français International de Kyoto, for her advice, her help in identifying several samples, and for her corrections to the initial version of my thesis. Her help was invaluable. I also extend my deepest gratitude to my father for accompanying me in the field, for cutting and polishing the samples, for the many discussions we had about this research and its execution, and for his assistance in improving this article.

Works Cited

Barnes G-L. (2003) Origins of the Japanese Islands: The New "Big Picture"., International Research Centre for Japanese Studies, National Institute for the Humanities, Japan Review, 15: 3-50. DOI: 10.15055/00000254.

Kameoka City History Compilation Committee (亀岡市史編さん委員会編) (1995) Overview of the Topography of Kameoka City and Its Surroundings (亀岡市およびその周辺の地形概観)., Newly Revised History of Kameoka City (新修 亀岡市史), 1(1-3). 872 p.

Ferry R.D., Chilingar, G.V. (1955) Summary of “Concerning Some Additional Aides in Studying Sedimentary Formations” by M.S. Shvetsov., Journal of Sedimentary Petrology, 25(3): 229-234. DOI: 10.1306/74d70466-2b21-11d7-8648000102c1865d.

Gianni M-G., Guo Z., F. Holt A., Faccenna C. (2025) Noncollisional orogeny in northeast Japan driven by nearby same-dip double subduction., Nature Geoscience, 18: 525-533. DOI: 10.1038/s41561-025-01704-5.

Hayase I., Tsutsumi T. (1956) The minute radioactive minerals of the Gyoja-yama granite, Kyoto Préfecture., Journal of the Japanese Association of Mineralogists, Petrologists and Economic Geologists, 40: 216-223. DOI: 10.2465/ganko1941.40.216.

Ishihara H., Ivanov V-V., Ratkin V. (1996) Gold Deposits and Their δ34S Ratios of the Sikhote-Alin, Russia – A Comparative Study with Those of the Sanyo Belt of Japan., Resource Geology, 46(5): 267-277. DOI: 10.11456/shigenchishitsu1992.46.267.

Isozaki Y., Yukio I. (2023) Ordovician Japan: geotectonic setting and palaeogeography., Geological Society of London, 533: 505-517. DOI: 10.1144/SP533-2022-80.

Kawabe T. (1989) Block movements of the Pliocene to Pleistocene sedimentary basins in Kinki district, Japan., Earth Science, 43(6): 401-416. DOI: 10.15080/agcjchikyukagaku.43.6_402.

Li Y-H., Ishihara S. (2013) Compositional variation of the late Cretaceous–Paleogene plutons from southwest Japan and its implication for ore genesis and continental growth., Journal of Asian Earth Sciences, 70–71: 142–159. DOI: 10.1016/j.jseaes.2013.03.009.

Lo P-C., Lo W., Wang T.T., Hsieh Y-C. (2021) Application of Geological Mapping Using Airborne-Based LiDAR DEM to Tunnel Engineering: Example of Dongao Tunnel in Northeastern Taiwan., Applied Sciences, 11(10). DOI: 10.3390/app11104404.

Nakajima T., Takahashi M., Imaoka T., Shimura T. (2016) Granitic rocks., The Geology of Japan: 251-272. DOI: 10.1144/GOJ.10.

Okada A., Kada A., Uemura Y., Togo M., Takemura K., Yoshioka R., Tsutsumi H., Umeda Y., Oika K., Matsui K., Sugimori T., Sugiyama N., Sonoda T., Umeda T., Matsumura N., Yamada K., Furusawa A. (2005) Quaternary faulting and subsurface structure of the Kameoka fault zone in Kyoto Prefecture, western Japan., Research committee for the active faults in Kyoto Prefecture. Active Fault Research, 25: 93-108. DOI: 10.11462/afr1985.2005.25_93.

Otofuji Y-I. (1996) Large Tectonic Movement of the Japan Arc in Late Cenozoic Times Inferred from Paleomagnetism: Review and Synthesis., The Island Arc, 5(3): 229-249. DOI: 10.1111/j.1440-1738.1996.tb00029.x.

Peacock D., Corke E. (2020) How to use a drone safely and effectively for geological studies., Geology Today, 36(4): 146-155. DOI: 10.1111/gto.12320.

Suga K., Yeh M-W. (2020) Secular Variation of Early Cretaceous Granitoids in Kyushu, SW Japan: The Role of Mélange Rocks as a Possible Magma Source., Frontiers in Earth Science, 8(95): 1-17. DOI: 10.3389/feart.2020.00095.

Taverna J. (1998) Modélisation mécanique du plissement de la lithosphère., Thèse de doctorat ; Spécialité : Géophysique – Géomécanique – Géochimie, Université de Grenoble I: 87 p.

Vansutre S., Ahmad F., Singh S. (2025) Drone: A Promising Tool In Applied Geology., Journal of Applied Geology and Geophysics, 13(5). DOI: 10.9790/0990-1305010109.

Wu T-J., Wu J., Okamoto K. (2022) Intra-oceanic arc accretion along Northeast Asia during Early Cretaceous provides a plate tectonic context for North China craton destruction., Earth-Science Reviews, 226(3): 1-15. DOI: 10.1016/j.earscirev.2022.103952,

Xu X., Huang Z., Jiang D., Zeng G., Dai L-Q. (2024) Remnants and fragments of the subducted paleo-Pacific plate: Constraints from geochemistry and geophysics., Science China, Earth Sciences, 67(10): 3041-3061. DOI: 10.1007/s11430-024-1352-y.

Anti-Knocking Agents in Gasoline: A Comprehensive Review of Mechanisms, Performance, and Environmental Impact By Dhairya Vij

Abstract

For over a century, a single combustion phenomenon has dictated how efficient a spark-ignition engine can be and what may be put in its fuel: knock. Knocking is what happens when the unburned fuel–air mixture ahead of the flame ignites on its own, all at once, instead of burning smoothly. Engine knock remains the central constraint on the efficiency and power density of spark-ignition (SI) engines, and the chemistry deployed to suppress it has shaped fuel policy for nearly a century. This review traces that history and its present tensions. It first examines tetraethyllead (TEL), the additive that defined twentieth-century fuel chemistry, evaluating the radical-scavenging mechanism behind its effectiveness and the public-health evidence that ultimately drove its global phase-out. It then re-examines the two laboratory metrics that continue to govern fuel specification—the Research Octane Number (RON) and the Motor Octane Number (MON)—and surveys experimental and historical evidence that the engine-weighting factor K relating these metrics to real knock behaviour has drifted toward zero and into negative territory as engine technology has evolved, leaving current standards misaligned with modern hardware. Finally, it appraises a spectrum of non-metallic and renewable alternatives—computationally screened aniline-like agents, petroleum-fraction blends, ethanol and acetone–butanol–ethanol (ABE) blends, and microbially derived biofuels—arguing that each successive generation of additive represents, in part, a response to the unintended consequences of the one before. Across these strands, the review identifies the modernisation of octane specification and the early integration of toxicological screening as the most pressing directions for future work.

Keywords spark-ignition engine; octane index; anti-knock agents; leaded gasoline phase-out; sustainable fuels.

Introduction

Why study the gasoline engine at all when the electric vehicle (EV) is widely presented as its successor? The question is a useful entry point, because the assumption embedded within it is incomplete. Despite the rapid growth of EVs, SI and other gasoline-fuelled engines continue to perform work that is not readily electrified. A great deal of machinery beyond the passenger car—equipment used in agriculture, construction, and a range of industrial settings—still depends on liquid fuels, which makes full electrification unrealistic in the near term. The transition to electric mobility is, moreover, a gradual process shaped by practical limits in charging infrastructure, the economics of battery manufacture, and broader technological readiness. Insufficient charging networks, constraints on battery production, high capital costs, and the difficulty of scaling EV technologies all act to slow the changeover, and alternative carriers such as hydrogen remain at an earlier stage of infrastructural and economic maturity still.

For these reasons, understanding how gasoline-based SI engines operate—and how they can be improved during a transition that will span decades—remains a live and consequential problem. Existing vehicle fleets, the long service lifetimes of gasoline-powered equipment, and uneven global access to charging infrastructure together guarantee the continued relevance of the SI engine throughout this period.

Engine knock

Knocking—also termed engine knock or detonation—is a critical abnormal-combustion phenomenon in SI engines and has been studied extensively because of its impact on engine performance, efficiency, and durability. It occurs when the unburned end-gas ahead of the flame front auto-ignites spontaneously, producing pressure oscillations that impose severe mechanical stress on engine components (Heywood; Wei et al.). Knock is governed by a combination of factors, including fuel octane rating, in-cylinder temperature and pressure, ignition timing, and mixture composition. Persistent knocking not only limits the allowable compression ratio (CR)—thereby constraining thermal efficiency—but also accelerates wear and raises the risk of catastrophic failure, which makes its mitigation a central concern of combustion research. Understanding the mechanisms and control strategies for knock is therefore essential to optimising SI engine design and improving fuel efficiency in modern internal-combustion engines (ICEs).

Octane rating and the Octane Index

Because octane rating is the fuel property most directly tied to knock resistance, it is worth establishing at the outset how it is measured. Both the RON and the MON are determined on the same type of test engine—a Cooperative Fuel Research (CFR) engine—but under markedly different conditions. RON is measured at lower speed with room-temperature intake air, a comparatively gentle protocol. MON raises the engine speed and pre-heats the intake mixture, producing a more thermally aggressive test. The difference between the two metrics defines the octane sensitivity, $S = RON - MON$, and this gap proves to matter a great deal in modern engines. To connect these laboratory numbers to real engine behaviour, researchers use the Octane Index, $OI = RON - K\cdot(RON - MON)$, equivalently written $OI = (1 - K)\cdot RON + K\cdot MON$ (Kalghatgi, “Part I”; Viayna et al. 1). The factor K is the key variable: it expresses how much weight a given engine places on RON relative to MON under its specific operating conditions. As the sections that follow demonstrate, K is far from constant—and its behaviour in modern hardware is the crux of the case for reforming octane specification.

Tetraethyllead (TEL): Mechanism, Effectiveness, and Phase-Out

Emissions from SI engines

Gasoline engines produce a broad range of harmful outputs: carbon monoxide (CO), unburned hydrocarbons (HC), nitrogen oxides (NOx), sulfur oxides (SOx), carbon dioxide

(CO2), and particulate matter (PM). Each carries its own consequences—some acute, some long-term—for human health and the environment. CO2, in particular, is the dominant greenhouse-gas concern, and the sheer scale of global gasoline combustion makes it a significant contributor to climate change. The engineering responses to these emissions fall into three broad categories: altering how the engine itself operates through combustion tuning and ignition strategy; treating the exhaust after combustion using catalytic converters and particulate filters; and modifying the fuel, whether through additives or alternative blends, to reduce harmful products at source. It is the third category to which anti-knock additives belong, and for most of the twentieth century the dominant such additive was TEL (Heywood).

Why TEL was so effective

TEL's dominance in twentieth-century fuel chemistry rested above all on efficiency. An extraordinarily small quantity of TEL delivered a large increase in octane number, making it the cheapest and most practical means of bringing knock under control at the CRs manufacturers required for competitive performance. Without it, the knock problem was genuinely difficult to solve by other means at the time. Alternative metallic additives such as methylcyclopentadienyl manganese tricarbonyl (MMT) or ferrocene existed, but none matched TEL's combination of effectiveness, ease of blending, and low cost (Viayna et al. 1). Its principal drawback—combustion by-products including lead oxides and lead halides—required chemicals to prevent deposit build-up, an inconvenience that, within the understanding of the era, was not regarded as a reason to abandon the additive (Ure et al.).

Chemical mechanism

Upon entering the combustion environment, TEL decomposes under heat and oxidises to fine particles of lead monoxide (PbO), and it is these particles, rather than the parent molecule, which prevent knocking. The experimental account shows that the PbO formed in the end gas destroys the peroxides and removes the radical intermediates that would otherwise drive the chain-branching reactions leading to auto-ignition (Salooja). By intercepting those species before the chain reaction reaches the self-ignition threshold, TEL confers additional stability on the fuel–air mixture at temperatures and pressures it would not otherwise tolerate, and the result is a higher measured octane number—greater resistance to knock relative to the iso-octane and n-heptane reference mixture. Because the lead species act in this catalytic, heterogeneous fashion, each particle consuming many radicals, TEL remains effective at molar dosages far smaller than any organic alternative can match, which is why the mechanism worked so well in practice (Salooja; Ure et al.).

Public-health evidence

That TEL carried serious costs is not a matter of inference but of a substantial and consistent evidence base built up over decades. Long before the additive was withdrawn, toxicologists and epidemiologists had documented its harms, and the body of work has only

grown more conclusive since: large systematic reviews now link leaded gasoline directly to elevated blood lead levels across populations, and trace its consequences from childhood neurological damage to cardiovascular disease in adults (Angrand et al.; Lacerda et al.). The strength of this evidence lies less in any single study than in its convergence across countries, decades, and methods, and it is this record that the following passages examine.

Blood lead levels and leaded gasoline

A systematic review (Angrand et al.) addressed a simple question: "When countries removed lead from gasoline, did population blood lead levels (BLLs) actually fall, and by how much?" A secondary objective was to quantify the ongoing public-health burden attributable both to TEL use and to the legacy contamination it left behind. Drawing on more than 38 independent studies across multiple continents, the answer was unambiguous. In every country examined, removing lead from gasoline was followed by significant declines in mean BLLs among both children and adults. In many settings BLLs continued to fall for years after leaded fuel had been eliminated entirely, reflecting the slow clearance of accumulated environmental contamination. The consistency of that pattern across such varied populations, climates, and healthcare systems is striking; it is not a single-country result or an artefact of a particular study design, but a global signal.

A similar study (Lacerda et al.) took a more quantitative and global view, tracking BLL trends across regions as TEL was reduced and eventually eliminated, assessing whether legacy contamination continued to affect populations, and comparing the pace and scale of recovery in different parts of the world. BLLs fell exponentially across continents as leaded gasoline was withdrawn. TEL in fuel was not merely a contributing factor to human lead exposure; it was the dominant one. The post-phase-out picture is more complicated: airborne lead dropped sharply, but legacy lead in urban soils continues to expose people, and children in particular, in cities where decades of combustion left deep deposits (Lacerda et al.).

The global phase-out

The removal of TEL from gasoline was not a quiet regulatory adjustment but a public-health intervention on a global scale, and it took decades to complete (Nriagu). For most of the twentieth century, TEL was simply part of how fuel worked: it boosted octane cheaply and reliably, it allowed manufacturers to build higher-compression engines, and the alternative—reformulating fuel from first principles—was considerably more expensive. The toxicological consequences of burning millions of tonnes of leaded fuel were either genuinely unknown in the early decades or, in some cases, known and quietly suppressed by commercial interests with a stake in keeping TEL on the market (Nriagu). As the evidence accumulated—brain damage in children, cardiovascular disease in adults, and persistent environmental contamination—the regulatory and scientific consensus slowly shifted (Lacerda et al.). Governments began to act from the 1970s onward, though the timeline varied enormously by country, with some regions completing the transition only in the early 2000s. Refineries

retooled, ethanol and other oxygenates replaced the lost octane, and the effect on public health was rapid and measurable (Angrand et al.). BLLs dropped dramatically, urban air quality improved, and the scale of the improvement confirmed what epidemiologists had argued for years (Angrand et al.; Lacerda et al.).

Two implications for future work follow. First, the persistence of legacy soil lead documented by Lacerda et al. indicates that remediation should be treated as an unfinished research and policy agenda. Second, and more generally, the bioaccumulative behaviour of lead establishes environmental persistence and clearance time as first-order screening criteria that any candidate replacement additive should be required to satisfy before large-scale deployment.

The Significance of RON and MON to Engine Performance

Having established the historical baseline, this section turns to the metrics that continue to govern fuel specification. The question is whether RON and MON, as currently measured and weighted, still describe the knock behaviour of modern engines—or whether the standards have been outpaced by the hardware. Two studies bear directly on this question: a focused experimental investigation of a single modern engine, and a historical analysis spanning seven decades of engine development. A third and more recent study (Ling et al.) extends the same question to a turbocharged direct-injection engine and to whole-vehicle drive-cycle fuel consumption, and is drawn upon below as corroborating evidence.

In most real gasolines RON and MON are chemically coupled, so their individual effects are hard to separate. Wang et al. investigated this phenomenon. They formulated six fuels specifically to decouple the two metrics, letting each be varied on its own while knock behaviour was tracked (Wang et al.). Standard CFR measurements were kept for RON and MON, but the point of the bespoke fuel set was to isolate which of the two actually does the work once knock becomes the limiting factor.

Experiments were conducted on an AVL single-cylinder direct-injection spark-ignition (DISI) research engine of 82 mm bore and 86 mm stroke, fitted with a direct piezoelectric injector, variable valve timing, and a central spark plug (Wang et al.). The six fuels were blended to hold other properties as constant as possible: heat of vaporisation, oxygen content, lower heating value (LHV), and stoichiometric air–fuel ratio (AFR) were all kept similar across the set. Testing covered stoichiometric operation at 1800 rpm, loads between 4 and 20 bar indicated mean effective pressure (IMEP), and three CRs of 9.5:1, 10.5:1, and 11.5:1. Each condition used the fuel-specific optimum spark timing, and both knock-limited spark advance (KLSA) and in-cylinder pressure traces were recorded throughout (Wang et al.).

When the engine was not knock-limited, RON and MON made almost no difference to indicated thermal efficiency (ITE)—as expected, since octane quality becomes relevant only once knock constrains the engine. The instructive results emerged under knock-limited conditions. Increasing MON did not improve ITE and, in some cases—particularly around RON 93 at high CR—actually degraded it. RON, by contrast, exerted a clear and consistent positive effect on efficiency (Wang et al.). The calculated K values for this engine were below 0.5 and

frequently negative, a significant result: it implies that a fuel with a larger RON–MON gap can outperform one with a higher MON in knock resistance. This cuts against the conventional assumption that MON, being the harsher test, is the more protective metric, and it is consistent with the broader octane-appetite literature in which high-sensitivity fuels resist knock in modern engines better than their MON alone would predict (Kalghatgi, "Part II"; Szybist and Splitter). The practical implication is that fuel specifications built around MON may be optimising for the wrong property in modern DISI engines, leaving efficiency unrealised.

What stands out in this dataset is the sign of K itself. Negative values overturn the hierarchy that specifications quietly assume, since they mean a high-sensitivity fuel can hold off knock better than a higher-MON one in a modern DISI engine (Wang et al.). That single result does a lot of work in the wider argument that octane standards need rethinking as engines change. It comes, though, from one research engine, so the obvious next step is to repeat the decoupled-fuel approach across a spread of production powertrains and see whether the distribution of K—rather than one striking value—holds up.

Rather than testing a single engine, a second study undertakes a historical review, combining archival vehicle data with targeted experiments across engines from different eras—from early carburetted designs to modern turbocharged direct-injection systems (Mittal and Heywood). This approach allows the authors to trace how decades of engineering advances—fuel injection, charge cooling, and boosted breathing—gradually shifted in-cylinder thermodynamic conditions away from those the CFR protocols were designed to replicate. The result is that K values have drifted toward zero and into negative territory over roughly seventy years of engine development (Mittal and Heywood). In practice this means RON has become the better predictor of knock resistance for modern engines, because such engines operate with lower charge temperatures at knock onset—conditions that resemble the RON test more closely than the MON test. The same fuel can behave differently across operating points within a single engine, let alone across engine generations. If regulators continue to mandate MON as a primary metric for engines in which K is already low or negative, the resulting fuel specifications will be systematically misaligned with actual engine needs, with downstream effects on formulation decisions, labelling standards such as the anti-knock index (AKI), and calibration strategies.

Convergent evidence and its limits

What gives these two studies their weight is that they reach the same place from opposite directions. Mittal and Heywood take the long view: across seventy years of engine development, they show that the gap between the test bench and the road is not the quirk of one modern design but a path the whole industry has followed while the standards stayed put. Wang et al. come at the same question from the laboratory, separating the two metrics in a single controlled engine and measuring the markedly negative K values that the historical trend had implied. A decades-long pattern and a careful single-engine measurement, neither relying on the other, end up agreeing—and it is that independence that makes the conclusion difficult to wave away. Their

shared lesson is worth stating plainly: a lone octane number says little about how a fuel will actually perform, because the same fuel behaves differently as conditions shift, and any case for reform ought to begin there.

Newer, vehicle-scale work points in the same direction. In a 2.3-litre turbocharged direct-injection engine, raising RON reliably improved both knock resistance and thermal efficiency under knock-limited operation, and the benefit carried through to lower fuel consumption across a full drive cycle, whereas MON made little consistent difference (Ling et al.). Set the three side by side—the historical sweep, the controlled experiment, and the production-vehicle test—and the case for overhauling octane specification reads as practical rather than academic. What stands in the way is largely procedural: a specification built on a measured K can become regulation only once there is an agreed, repeatable way of determining K across the full range of engines on the road, and settling on that method is the work that should come first.

Sustainable and Non-Metallic Anti-Knock Additives

If MON-based specification is increasingly obsolete, the question becomes what should replace TEL. Recent reviews survey a broad landscape of candidate octane enhancers—bioethanol, prenol, furan mixtures, di-isobutylene, and related oxygenates—and trace the market and quality requirements now driving the search for environmentally acceptable high-octane gasoline (Abdellatief et al., “Recent Trends”; Abdellatief et al., “Advanced Progress”). Building on that landscape, this section surveys three lines of response—computational screening of organic agents, alcohol and bio-alcohol blends, and microbially derived fuels—and closes with a comparative assessment of ethanol against the TEL baseline established earlier.

Computational screening of aniline-like agents

Traditional additive discovery leans heavily on trial-and-error engine testing, which is slow and expensive. Viayna et al. set out to short-circuit that, starting from a very large chemical space—several million aniline-like molecules—and narrowing it through a tiered computational pipeline. An initial filter applies practical blending criteria (molecular weight, aqueous solubility, similarity to aniline) to cut the field to a workable subset; density functional theory (DFT) calculations then estimate how well each candidate stabilises the radical intermediates that drive knock. Only the strongest computational performers are synthesised, blended into a reference gasoline, and measured for RON on a standard engine. It is this order—screen first, synthesise selectively, test last—that sets the work apart (Viayna et al. 2–4).

Several of the aniline-derived compounds performed well, lifting RON appreciably at low blend concentrations and, in a few cases, beating conventional alcohols or ethers at the same or lower dosages (Viayna et al. 4–5). A well-chosen architecture appears to do more work per unit added than a larger dose of something less carefully designed.

The practical upshot is that the pipeline may matter more than the particular molecules it happened to find. Front-loading discovery with computation weeds out weak candidates before anyone runs an engine, which saves a good deal of wasted effort, and—more importantly—it lets environmental and health screening sit at the start of the process rather than being bolted on at the end. Given the TEL precedent, that ordering is not just tidy but a direct lesson from history. One limitation worth flagging is that efficacy and toxicity are still handled as separate filters; coupling the DFT descriptors used for performance with quantitative toxicity and persistence measures, so a single stage ranks candidates on both at once, would be a natural improvement.

Ethanol–gasoline blends

Ethanol is, in practice, the oxygenate that took over much of the octane duty once lead was removed, so how it behaves in a running engine is more than an academic question. Szabó et al. look at what happens when 10% and 30% of the gasoline is swapped for ethanol in a working SI engine, and they are careful to ask not just whether efficiency moves but why—tracing the effect back to ethanol's high octane number, the charge-cooling from its large latent heat of vaporisation, and the extra oxygen it brings to the mixture. Standard gasoline is the baseline, and the measured outputs run from brake-specific fuel consumption (BSFC) and thermal efficiency to combustion behaviour and exhaust composition.

Ethanol blends resist knock better than neat gasoline, which the authors put down to ethanol's naturally high octane number—the property that lets an engine run higher compression and reap the efficiency that follows. The added oxygen pushes combustion toward completion, which can trim some emissions and nudge power up, while the charge-cooling from ethanol's latent heat improves how much mixture the cylinder draws in, though it can also make cold starts harder. Push the ethanol fraction very high and the lower energy density starts to bite, so fuel economy slips a little. Within the range tested, the sweet spot appears to sit around E10 to E30, where the overall efficiency picture comes out positive (Szabó et al.).

Acetone–butanol–ethanol (ABE) blends

Tang et al. extend the bio-alcohol blending account beyond ethanol to a three-component mixture: acetone, butanol, and ethanol in the ABE ratio produced by fermentation. Blends of 10%, 20%, and 30% ABE were tested in a high-speed SI engine, with knock intensity tracked through a calculated knocking factor. Across all blend ratios, ABE reduced knock tendency relative to pure gasoline, and the improvement scaled with concentration. The oxygenated nature and high octane rating of the alcohol components are credited with stabilising combustion and pushing back the auto-ignition threshold of the fuel–air mixture. The study's scope is intentionally narrow—ABE versus gasoline and nothing else—so it does not position itself within a broader additive comparison; but as a focused experimental dataset on how a renewable fermentation-derived blend behaves at the knock threshold, it constitutes directly useful evidence.

The steady, concentration-scaled improvement is itself a useful result: within the range tested, knock suppression rose predictably with the blended fraction, which makes the behaviour straightforward to apply in practice. What ABE adds beyond that is reach—it widens the renewable end of the additive space and shows that a fermentation-derived blend can compete on anti-knock performance while coming from a renewable route. An open question is how it stacks up against ethanol head-to-head at matched octane index, judged not only on knock but on the full emissions and energy-density picture, so that the renewable advantage can be weighed against any combustion penalty.

Comparative assessment: ethanol versus TEL

The contrast between TEL and ethanol as anti-knock agents is stark, and it extends well beyond the lead-versus-no-lead distinction. TEL combustion releases lead oxides and lead halides—persistent, non-biodegradable compounds that accumulate in soil, water, and biological tissue—whereas ethanol produces none of these. Lead bioaccumulates through the food chain and causes neurological and systemic damage; ethanol is comparatively benign and does not persist in the environment. The environmental legacy of TEL extended decades beyond its use, as documented earlier, while ethanol breaks down naturally (Lacerda et al.). On the production side, TEL is derived from petrochemical feedstocks, whereas ethanol can be produced from sugarcane, corn, or other renewable crops. Ethanol combustion also tends to yield lower CO and fewer unburned HC than TEL-laced fuel.

The significant caveat is that large-scale ethanol production creates its own environmental pressures: competition with food crops for land, deforestation risk in some regions, and high water consumption (Searchinger et al.). The comparison is therefore considerably better than it appears on the TEL side, but not quite as unambiguously clean as it appears on the ethanol side. This asymmetry is precisely what motivates the search for production routes that avoid both the toxicity of TEL and the land-use burden of first-generation ethanol—the subject of the final section.

So ethanol is not a clean swap for TEL: the performance and emissions gains are real, but so are the trade-offs in energy density and cold starting, and both matter in service. The point that moderate blends beat both neat gasoline and very high ethanol ratios is arguably the most policy-relevant one, and it would be worth pinning down the optimal blend window as a function of engine K value. That loops back to the earlier question of whether octane quality should be specified through an engine-aware Octane Index built on a measured K rather than through fixed RON and MON (Kalghatgi, “Part I”; Mittal and Heywood).

Microbially derived biofuels

With greenhouse-gas pressure mounting and fossil reserves finite, an octane booster that a microorganism can make is an appealing prospect, and that is the ground Mack et al. explore. They rate four candidate compounds from microbial metabolism—the isopentenols 3-methyl-2-buten-1-ol and 3-methyl-3-buten-1-ol, the branched alcohol isobutanol, and the

cyclic terpene limonene—by RON and by blending RON (BRON) on a single-cylinder CFR engine, against primary reference fuels (PRFs). Three of the four came out strong: the isopentenols and isobutanol returned high BRON values (roughly 118 to 126), a clear improvement in anti-knock quality, while limonene was only modest at a BRON near 99. What makes the result interesting is provenance—each of these can in principle be grown through a microbial pathway rather than synthesised from petroleum, which puts them forward as renewable, biomass-derived routes to octane enhancement.

Read as the experimental anchor of a broader third-generation argument, the study earns its place. TEL carries a toxicological cost and first-generation ethanol carries a land-use one; biofuels from microbial metabolism hold out the prospect of dodging both, since the feedstock need not compete with food and the products lack lead's persistence. The high BRON figures establish that the chemistry and combustion are promising. What remains unclear—and this is the real gap—is whether such pathways can hit the titres, yields, and costs that fuel-scale production demands. Later work has begun to push in this direction, toward microbially produced high-octane olefinic esters with good octane sensitivity, but a standardised, independent life-cycle assessment of carbon balance and biodegradability is still the precondition for calling these agents genuinely sustainable rather than merely promising (Carruthers et al.).

Conclusion

Looking across the studies gathered here, one pattern is hard to miss: each generation of anti-knock chemistry solved the pressing problem of its day while seeding the one that came next. TEL fixed knock cost effectively and well, and in doing so imposed a public-health cost that took decades to recognise and longer to undo; its phase-out was a genuine policy success but left soil contamination that is still with us. The metrics carried over from that era—RON and MON—still govern specification even though the engine-weighting factor K has slipped toward and below zero, which leaves MON-anchored standards increasingly out of step with the hardware they describe. The most concrete conclusion to draw is that octane specification ought to move toward engine-aware indices built on a measured K, once a reproducible way to measure it exists.

On the additive side, the arc runs from computationally screened organic agents and simple petroleum-fraction blends, through ethanol and ABE blends that deliver real but qualified benefits, to microbially derived biofuels that aim to escape both the toxicity of TEL and the land-use burden of first-generation ethanol. None is without trade-offs—that is the recurring lesson—and choosing an additive is always a balance of performance, practicality, and environmental cost. The thread tying the sections together is that each generation answers the unintended consequences of the last, which is also why the two priorities that keep surfacing—modernising octane specification and moving toxicological screening upstream—matter more than any single new molecule.

In the future, progress will depend less on discovering new effects than on agreeing reproducible procedures—for measuring K, for ranking candidate additives (Viayna et al. 7), and

for accounting honestly for the full life cycle of a fuel (Abdul-Manan and Kalghatgi)—so that what is already understood in the laboratory can be carried into specification and policy. No anti-knock agent is without cost, from the toxicity of lead to the land-use burden of first-generation ethanol; the value of the work ahead lies in specifying, screening, and assessing fuels with the foresight that earlier decisions lacked.

Works Cited

Abdellatief, Tamer M. M., et al. “Advanced Progress and Prospects for Producing High-Octane Gasoline Fuel toward Market Development: State-of-the-Art and Outlook.” Energy & Fuels, vol. 37, no. 23, 2023, pp. 18266–18290. pubs.acs.org/doi/abs/10.1021/acs.energyfuels.3c02541

“Recent Trends for Introducing Promising Fuel Components to Enhance the Anti-Knock Quality of Gasoline: A Systematic Review.” Fuel, vol. 291, 2021, article 120112. sciencedirect.com/science/article/abs/pii/S0016236120331094

Abdul-Manan, Amir F. N., and Gautam T. Kalghatgi. “Exploring Alternative Octane Specification Methods for Improved Gasoline Knock Resistance in Spark-Ignition Engines.” Frontiers in Mechanical Engineering, vol. 4, 2018, article 20. https://doi.org/10.3389/fmech.2018.00020

Angrand, Rachel C., et al. “Relation of Blood Lead Levels and Lead in Gasoline: An Updated Systematic Review.” Environmental Health, vol. 21, 2022, article 138. https://doi.org/10.1186/s12940-022-00936-x

Carruthers, David N., et al. “Microbial Production of High Octane and High Sensitivity Olefinic Ester Biofuels.” Biotechnology for Biofuels and Bioproducts, vol. 16, 2023, article 60. https://doi.org/10.1186/s13068-023-02301-7

Heywood, John B. Internal Combustion Engine Fundamentals. McGraw-Hill, 1988. Google Books

Kalghatgi, Gautam T. “Fuel Anti-Knock Quality—Part I: Engine Studies.” SAE Technical Paper 2001-01-3584, 2001. https://doi.org/10.4271/2001-01-3584

---. “Fuel Anti-Knock Quality—Part II: Vehicle Studies—How Relevant Is Motor Octane Number (MON) in Modern Engines?” SAE Technical Paper 2001-01-3585, 2001. https://doi.org/10.4271/2001-01-3585

Lacerda, et al. “Global Decrease in Blood Lead Concentrations Due to the Removal of Leaded Gasoline.” Chemosphere, 2023. sciencedirect.com/science/article/abs/pii/S0045653523004745

Ling, Zhe, et al. “Impact of Octane Numbers on Combustion Performance and Driving Cycle Fuel Consumption of a Turbocharged Direct-Injection Spark-Ignition Engine.” Fuel, 2024. sciencedirect.com/science/article/pii/S0016236124014352

Mack, J. Hunter, et al. “Investigation of Biofuels from Microorganism Metabolism for Use as Anti-Knock Additives.” Fuel, vol. 117, part B, 2014, pp. 939–943. sciencedirect.com/science/article/abs/pii/S0016236113009654

Mittal, Vikram, and John B. Heywood. “The Shift in Relevance of Fuel RON and MON to Knock Onset in Modern SI Engines over the Last 70 Years.” SAE International Journal of Engines, vol. 2, no. 2, 2010, pp. 1–10. https://doi.org/10.4271/2009-01-2622

Nriagu, Jerome O. “The Rise and Fall of Leaded Gasoline.” Science of the Total Environment, vol. 92, 1990, pp. 13–28. https://doi.org/10.1016/0048-9697(90)90318-O

Salooja, K. C. “Studies Relating to the Mechanism of Antiknock Action of Tetraethyl Lead.” Combustion and Flame, vol. 9, no. 3, 1965, pp. 211–217. sciencedirect.com/science/article/abs/pii/0010218065900878

Searchinger, Timothy, et al. “Use of U.S. Croplands for Biofuels Increases Greenhouse Gases through Emissions from Land-Use Change.” Science, vol. 319, no. 5867, 2008, pp. 1238–1240. https://doi.org/10.1126/science.1151861

Szabó, et al. “Comprehensive Efficiency Analysis of Ethanol–Gasoline Blends in Spark-Ignition Engines.” Eng, 2025. mdpi.com/2673-4117/6/10/256

Szybist, James P., and Derek A. Splitter. “Pressure and Temperature Effects on Fuels with Varying Octane Sensitivity at High Load in SI Engines.” Combustion and Flame, vol. 177, 2017, pp. 49–66. https://doi.org/10.1016/j.combustflame.2016.12.002

Tang, Qijun, et al. “Effect of Acetone–Butanol–Ethanol (ABE) and Gasoline Blends on Knocking Performance.” Thermal Science and Engineering Progress, vol. 46, 2023, article 102175. sciencedirect.com/science/article/abs/pii/S2451904923005280

Ure, Andrew D., et al. “Rational Design and Testing of Anti-Knock Additives.” Energies, vol. 13, no. 18, 2020, article 4923. https://doi.org/10.3390/en13184923

Viayna, Antonio, et al. “Holistic Approach to Anti-Knock Agents: A High-Throughput Screening of Aniline-Like Compounds.” Fuel, vol. 305, 2021, article 121518. https://doi.org/10.1016/j.fuel.2021.121518

Wang, Chongming, et al. “Significance of RON and MON to a Modern DISI Engine.” Fuel, vol. 209, 2017, pp. 172–183. https://doi.org/10.1016/j.fuel.2017.07.071

Wei, Haiqiao, et al. “Effect of Flame Speed on Knocking Characteristics for SI Engine under Critical Knocking Conditions.” Fuel, 2020. sciencedirect.com/science/article/abs/pii/S0016236120318421

The Effects of pH Level on the Decomposition of Xenobiotic Vinyl Compounds and the Resulting Health and Environmental Issues By Stella Dybczakova

I. Abstract

Vinyl compounds have a significant impact in our daily lives and the environment. Examples of commonly used vinyl compounds include: PVC (polyvinyl chloride) and EVA (ethylene vinyl acetate) [1]. They can be found frequently in products used in construction, flooring and packaging due to their long-lasting durability. However, these materials are also xenobiotic, meaning they are produced artificially and cannot be degraded in biological systems [2]. This paper discusses the influence of pH on vinyl compounds and how it affects their degradation processes, which releases harmful chemicals affecting our environment and health. The pH of a substance can alter the quality and efficiency of the decomposition. Understanding and researching the behavior of pH can help us comprehend how these vinyl compounds affect our health and environment, while also developing more effective protective and efficient decomposition procedures.

II. Introduction

The widespread use of vinyl compounds has raised concerns about their environmental and health impacts. We encounter vinyl compounds everyday, from drinking water which flows through pipes to walking on floors both principally made of vinyl chlorides [3]. Unlike natural polymers, which can easily break down in nature, vinyl compounds cannot be broken down due to their xenobiotic origins. This means that once they are released into the environment, these materials accumulate in natural habitats such as oceans and soils, polluting the area and posing severe risks to animals and humans [2].

Another crucial factor affecting the decomposition of vinyl compounds is the pH of the surrounding environment. Environmental pH is never constant, with acidic ocean zones created by the accumulation of carbon dioxide, or basic volcanic areas [4]. These contrasts help researchers find new ways of recycling and fully decomposing vinyl compounds. Furthermore, the level of pH can release additives from these compounds, such as plasticizers, that can be harmful to our health and the surrounding environment. pH not only helps in the degradation of vinyl compounds, but can also harm living organisms by releasing harmful substances into the air [5].

The influence of pH on vinyl compounds is essential for learning more about their risks, designing safer, and more biodegradable products.

III. Chemistry and Environmental Effects of Vinyl Compounds

The vinyl group is a functional group in organic chemistry; meaning that every compound in the group has a similar structure of (-CH=CH2) [6,7]. The vinyl group is a type of the alkenyl group, characterized by the two carbon atoms sharing a double-bond as depicted in

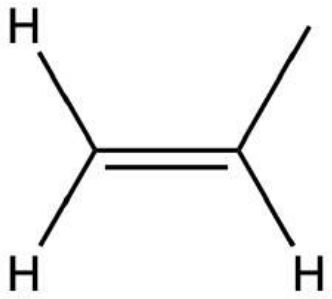


Figure 1, basic structure of a vinyl compound

Figure 1. The first carbon atom has a dash in front of it, this means that this carbon atom is bonded to another atom within the molecule [8].

The reason behind why vinyl compounds are so hard to break down is attributed to the strong double-bond between the two carbon atoms.

However, during the polymerization of vinyl compounds, that is Started by heat or light begins the process of this carbon double-bond opening up and forming long chains, this is depicted in Figure 2 [9]. This opening up causes the double-bond to change from reactive to stable and single bonded. This structural change is the reason why vinyl compounds are so durable. These bonds are capable of surviving natural decomposition processes like hydrolysis, meaning having a reaction with water, leaving them to stay in the environment unable to decompose for centuries or even a millennium [10]. The addition of plasticizers, which were originally meant to enhance the flexibility, ends up releasing toxic chemicals over time, raising concerns about the environmental and health effects [3].

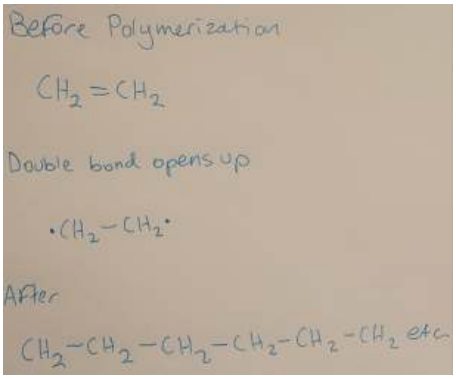


Figure 2, polymerization of a vinyl compound

Over time, UV from the sun and waves (if in the ocean) start to break down these products to create microplastics. These tiny, microscopic pieces of plastic enter our rivers, soils and have even been found in the air we breathe [11]. Every year, the growth of vinyl compounds in our environment grows due to its common use and resistance to break down. The persistence of these harmful compounds in the environment has created distress about our future and adjustment to life.

IV. Impact of pH on Vinyl Compounds Decomposition

The pH of a substance is central to decomposing of persistent vinyl compounds that remain in our environments for millennia. The three levels of pH each have distinct systems. Grasping these different ways is key to coming up with new ways to resolve the large amount of them in our ecosystems.

In neutral environments, which is any pH from around 7 to 8, vinyl compounds are broken down by bacterial enzymes. However, this process of bacterial and enzyme decomposition can take thousands of years to finish with optimal conditions [12]. Some other ways that can aid the degradation at neutral pH is the use of photochemical processes and the burning of the products, however these processes can take from a few hours to thousands of

years [13]. The decomposition of vinyl compounds at neutral pH is a very complex and long process.

In acidic environments, which is any pH from 0 to 6, vinyl compounds tend to degrade faster. For the most common kind of vinyl compound, PVC, the accelerated process is called "dehydrochlorination". Dehydrochlorination is an autocatalytic reaction, meaning that the reaction can produce and speed up its own rate of reaction. Dehydrochlorination involves the release of the HCl (hydrogen chloride) of the (Cl-CH=CH2) polymer chain. This initial elimination of HCl commences the autocatalytic decomposition of the rest of the chain [14, 15]. The decomposition also leads to releasing of toxic additives and metal stabilizers, harming the exposed population. In the end, the decay of vinyl compounds, primarily PVC, in acidic pH leads to not just a sped up reaction but also the release of harmful substances.

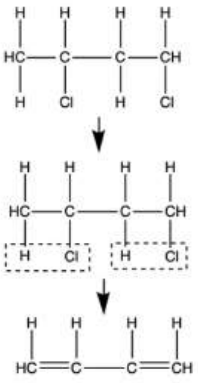


Figure 3, process of dehydrochlorination

In basic environments, which is any pH above 8, the decomposition of vinyl compounds is much slower than acidic and neutral pH. This is because most vinyl compounds are more stable in basic conditions than in acidic and neutral environments. However, there are some exceptions to this stability, similar to neutral pH, some organisms, such as halophiles thrive more in basic environments aiding in the decay of these compounds [16]. Another method of decomposing vinyl compounds in basic pH is by "cracking". An experiment conducted by Bishop et al. found that high pH leads to "stressful" conditions, causing the cracking, or breakdown of the PVC structure. Meaning that the physical conditions of the PVC becomes more brittle and decomposes easier. But, cracking is still not as reliable as the dehydrochlorination process found in lower pH's [17]. In conclusion, decomposing vinyl compounds in basic pH environments is generally inefficient due to the slower reaction rates and the lack of key mechanisms such as dehydrochlorination, which significantly accelerates the degradation in acidic conditions. Therefore, decomposition of vinyl compounds in basic pH is susceptible to microbial degradation, which tends to be much slower and weaker than the processes used in acidic conditions [18].

In summary, the pH factor plays a large role in the degradation of vinyl compounds such as PVC. The most reliable pH for degradation of vinyl compounds is acidic, meaning pH from 0 to 6. Acidic conditions use accelerating mechanisms, including dehydrochlorination, where the released hydrogen chloride from the structure acts as an autocatalyst, starting the quick breakdown of the polymer chain. In contrast, neutral and basic environments, which are characterized by pH greater than 7, have a very slow microbial degradation process which can last thousands of years to be completed. Comprehending the importance of pH in degradation can be pivotal in finding new ways to recycle and manage our use of xenobiotic vinyl compounds.

V. Health and Environmental Effects of Vinyl Compounds

Vinyl compounds, particularly vinyl chloride, present severe health risks to humans. Short-term exposure to these harmful compounds can result in a variety of symptoms such as dizziness, coughing, eye irritation, headache and central system depression [19]. Visible skin symptoms from physical contact with vinyl chloride include, skin rashes and frostbite due to the rapid evaporative cooling, which can develop into blisters and peeling [20]. Individuals who suffer with chronic illnesses such as asthma may promote more rapid worsening of symptoms. Yet short-term symptoms usually clear up after a few hours or days.

Individuals who are exposed for a longer and consistent period of time to vinyl chloride can experience more severe symptoms. Workers who are regularly working with the compound can develop chronic diseases linked to vinyl chloride. Over time, exposure causes damage to organs principally the liver, the conditions range from enlargement of the organ to a rare aggressive liver cancer called hepatic angiosarcoma [21]. For these reasons, U.S. major health agencies, including the National Cancer Institute, the U.S. Environmental Protection Agency and the National Institutes of Health, recognize vinyl chloride as a carcinogen.

PVC and other common vinyl compounds have been found in numerous waste sites all over the United States. Exposure to PVC frequently occurs through respiration of contaminated air. There are over 373,000 U.S. citizens living in a three-mile radius of these waste sites. These sites are located in states such as Texas, Louisiana, Kentucky, New Jersey, releasing over 414,000 pounds of carcinogenic vinyl chloride in 2021 [19,22]. All of these individuals are put in severe risk of being exposed to harmful levels of vinyl chloride and other compounds. Since vinyl chloride can be transmitted through air, it can travel further from the initial site, meaning that it can go into our local water sources and soils. This continued vulnerability, unpreventably affects sensitive groups of people such as children and the elderly, causing a large spike in the number of chronic diseases.

Despite a great deal of scientific evidence showing the serious consequences which vinyl compounds have on the health and environment, many are still unaware. Individuals who live near manufacturing and waste sites are oblivious of the health issues that vinyl compounds cause. Vulnerable populations, made up of children and elderly, endanger their well-being constantly as a result of not being educated about the consequences of exposure to vinyl compounds.

Solutions for finding more profitable and effective ways to stop polluting and harming populations are lacking attention. However, there are ideas such as using green chemistry to find alternative, bio-vinyl compounds [23] and using the previously discussed pH for easier degradation and less contamination.

In the end, by increasing public awareness and understanding the health and environmental impacts of vinyl compounds we could reduce exposure, protect ecosystems and ensure that our habitats aren't polluted.

VI. Conclusion

In summary, vinyl compounds, especially vinyl chloride, are present in our day-to-day lives. However, this persistence of products in our lifetimes causes serious consequences for our health and environment. Exposure can range from a spectrum of negative effects. From the short-term headaches and dizziness, to the long-term illnesses, including cancer. This underscores the importance of finding an effective way to break down these vinyl compounds. The slow and lengthy process of degrading these compounds under neutral pH is a clear example of why these compounds stay so long in our ecosystems. Their breakdown can be accelerated in acidic substances through the auto-catalytic process of dehydrochlorination. Nevertheless, this process is still being researched due to the toxic byproducts. These concerns highlight the urgent need to regulate and find new ways for finding an efficient degradation, ensuring a safer future for humanity and our planet.

Works Cited

1.Britannica Editors. “Vinyl Compound | Synthesis, Polymers, Monomers.” Encyclopedia Britannica, 20 July 1998, www.britannica.com/science/vinyl-compound.

2.Garg, Manisha. "Xenobiotic Compounds: Meaning, Hazards and Biodegradation." Biology Discussion, 16 Sept. 2016, www.biologydiscussion.com/microbiology-2/bioremediation/xenobiotic-compounds-meaning-hazards-and-biodegradation/55625.

3.SpecialChem. Polyvinyl Chloride (PVC): How to Select the Right Grade? 21 Apr. 2026, www.specialchem.com/plastics/guide/polyvinyl-chloride-pvc-plastic.

4.Extreme Habitats Around the Globe. education.nationalgeographic.org/resource/extreme-habitats-around-globe/3rd-grade.

5.Agency for Toxic Substances and Disease Registry, et al. "Potential for Human Exposure." Vinyl Chloride, 2023, www.atsdr.cdc.gov/toxprofiles/tp20-c5.pdf.

6."What Does Vinyl Mean in Organic Chemistry?" YouTube, uploaded by Chemistry For Everyone, 13 Jan. 2025, www.youtube.com/watch?v=px3NY7noEdg.

7.Bull, C. "Vinyl Compounds." Chemical Bull Pvt Ltd, www.chemicalbull.com/chemicals/vinyl-compounds.

8.The Editors of Encyclopaedia Britannica. "Vinyl Compound: Synthesis, Polymers, Monomers." Encyclopedia Britannica, 8 Aug. 2025, www.britannica.com/science/polyvinyl-chloride.

9.Singha, Subrata, et al. "Recent Developments on Cationic Polymerization of Vinyl Ethers." ACS Polymers Au, vol. 4, no. 3, 2024, pp. 189–207, https://doi.org/10.1021/acspolymersau.3c00055.

10.Tebes-Stevens, Caroline, et al. "Prediction of Hydrolysis Products of Organic Chemicals under Environmental pH Conditions." Environmental Science & Technology, vol. 51, no. 9, 2017, pp. 5008–16, https://doi.org/10.1021/acs.est.6b05412.

11.Campisi, L., et al. "Polyvinyl Chloride (PVC), Its Additives, Microplastic and Human Health: Unresolved and Emerging Issues." The Science of the Total Environment, vol. 960, 2025, p. 178276, https://doi.org/10.1016/j.scitotenv.2024.178276.

12.Bahl, S., et al. "Biodegradation of Plastics: A State of the Art Review." Materials Today: Proceedings, vol. 39, 2020, pp. 31–34, https://doi.org/10.1016/j.matpr.2020.06.096.

13."How Long Does It Take for Plastic to Decompose?" Chariot Energy, 3 June 2025, chariotenergy.com/blog/how-long-until-plastic-decomposes/.

14."Why Do We Need Stabilizers? PVC Degradation Mechanism." Westlake Akishima, 12 May 2022, www.westlake-akishima.com/2022/05/12/pvc-degradation-and-stabilization/?lang=en.

15.Hanopolskyi, A., et al. Autocatalysis: Kinetics, Mechanisms and Design, 2020, bioinf.boku.ac.at/pub/793.301/2020w/Flamm/hanopolskyi-CSC-rev-2020.pdf.

16.Chia, X. K., et al. "Role of Extremophiles in Biodegradation of Emerging Pollutants." Topics in Catalysis, 2024, https://doi.org/10.1007/s11244-024-01919-7.

17.Bishop, S., et al. "Environmental Stress Cracking of Poly(vinyl Chloride) in Alkaline Solutions." Polymer Degradation and Stability, vol. 70, no. 3, 2000, pp. 477–84, https://doi.org/10.1016/s0141-3910(00)00144-0.
18.Khan, M. F., et al. "Recent Advances in Fungal Xenobiotic Metabolism: Enzymes and Applications." *World Journal of Microbiology and Biotechnology*, vol. 39, no. 11, 2023, https://doi.org/10.1007/s11274-023-03737-7.
19.Agency for Toxic Substances and Disease Registry, and Case Studies in Environmental Medicine. Vinyl Chloride, www.19january2017snapshot.epa.gov/sites/production/files/2016-09/documents/vinyl-chloride.pdf. Accessed 17 Aug. 2026.
20.Goodman, Range S., et al. "Public Health Risks, Dermatological Manifestations, and Environmental Justice Associated with Vinyl Chloride Exposure: Narrative Review." JMIR Dermatology, vol. 6, 2023, p. e48998, https://doi.org/10.2196/48998.
21."Angiosarcoma of the Liver Among Polyvinyl Chloride Workers — Kentucky." Centers for Disease Control and Prevention, 3 Mar. 1999, www.cdc.gov/mmwr/preview/mmwrhtml/lmrk103.htm.
22."PVC Plastic and Environmental Justice." Toxic-Free Future, 13 Apr. 2023, toxicfreefuture.org/research/pvc-poison-plastic/pvc-plastic-and-environmental-justice/.
23."Bio-Vinyl Compounds: A Sustainable Leap for Vinyl Manufacturing." United Record Pressing, 6 Feb. 2025, www.urpressing.com/client-resources/blogs/bio-vinyl-compounds-for-sustainable-vinyl-manufacturing/.

Running on Empty?: Galactic Quenching Mechanisms and Gas Retention through Morphology, Mass, and Environment By Aarav Verma

Abstract

Over time, galaxies approach a stage of quiescence, meaning they have significantly halted star formation. Galaxy quenching has traditionally been modeled as a permanent process that requires all star-forming material to be expelled or consumed. However, recent observations have revealed persisting gas reservoirs in quiescent galaxies, challenging previously known quenching mechanisms. This study examines mechanisms that suppress star-formation across galaxies varying by mass, shape, and environment. Morphological transformations stabilize the gas disc against gravitational collapse, preventing the formation of stars through a process named morphological quenching. This process allowed for the presence of molecular gas within a galaxy without star formation. In interacting galaxies, rather than morphological quenching, interactions heated the gas which prevented it from collapsing. While higher-mass galaxies experienced long-term shutdown, some lower-mass galaxies experienced quenching cycles depending on their environment. Lower-mass galaxies were able to replenish their reservoirs of cold molecular gas from the surroundings and eject the warm gas, allowing for rejuvenated star formation. These results point to galactic quenching as a multi-step process. Fully understanding quiescent galaxies will require further studies with higher resolution simulations in the future.

Introduction

Star formation in galaxies is measured by the specific star formation rate (sSFR), which determines if a galaxy can be considered quenched, meaning it has stopped effectively making stars. The threshold to classify a galaxy as quiescent is when the time it takes for the galaxy to double its mass is at least five times longer than the age of the Universe at the epoch it is being observed.

The Star Formation Main Sequence is a correlation between stellar mass and star formation rate. When galaxies stop star formation and transition into early-type systems, their SFR drops much lower than the main sequence, falling into quiescence.

Galaxies are often categorized by being early-type or late-type. Early-type galaxies often are elliptical and have smooth and round stellar distributions without spiral arms. Usually, early-type galaxies are dominated by old and red stellar populations with low gas content. Late-type galaxies include spiral and irregular galaxies and are abundant in molecular gas and dust with a younger population.

The majority of the data on high-redshift quiescent galaxies has come from the past decade through the establishment of the James Webb Space Telescope (JWST). Early cosmological models such as EAGLE and SIMBA predicted a very small amount of quiescent galaxies in the early universe, believing that galaxy quenching was a slow, gradual process (Baker et al.). However, recent large-scale surveys across the redshift of z = 2-7 reveal a much greater amount of massive, quenched galaxies compared to simulation predictions (Baker et al.).

The shutdown of star formation was thought to be caused by internal processes such as active galactic nuclei (AGN) feedback, where supermassive black holes deplete molecular gas to inhibit star formation. However, with the discovery of more massive quiescent galaxies in the early universe, the belief shifted into the depletion of molecular gas due to a period of starburst, which is a highly accelerated period of star formation.

For example, a study done by de Graaf et al. discovered the massive quiescent galaxy RUBIES-EGS-QG-1 at the redshift of $z = 4.9$, which is about 1.2 billion years after the Big Bang. The galaxy accumulated a stellar population of 10^{11} solar masses during a period of only about 200 million years. After the starburst, star formation dropped rapidly until the galaxy quenched. Cosmological models predicted that the extreme starburst and early death of the galaxy was too rare to occur to be discovered by JWST. The discovery of it proved that quenching either had different causes, occurred much earlier, or was much more efficient in exhausting star-forming material.

Further discoveries found that quenching of ultramassive galaxies at early redshifts can be accelerated by environmental factors as well. At $3 < z < 5$, ultramassive quiescent galaxies form in cosmic overdensities, which are the densest regions in the cosmic web (Jespersen et al.). The higher mass concentration in the overdensities generates a stronger gravitational pull on cold gas. This accelerates the gas and allows it to funnel into the central galaxy, fueling a rapid and early period of starburst.

A study conducted by Skarbinski et al. found further evidence of starburst-driven quenching pathways, which was widespread during Cosmic Noon, the peak period where stars and galaxies formed at the fastest rate. The article studied 24 massive quiescent galaxies at $1 < z < 3$ and found that the star formation histories of these galaxies included a period of massive starburst before dropping. However, the timeline of quenching for these stars ranged from 0.06 to 1.75 gigayears (Gyr), meaning multiple quenching pathways occurred (Skarbinski et al.).

Further data from JWST supports the rapid quenching of galaxies. Galaxy GS-9209 was spectroscopically confirmed at $z = 4.568$ (Carnall et al.). Similarly to RUBIES-EGS-QG-1, the galaxy formed its stellar mass of 3.8×10^{10} solar masses in about 200 million years. However, this galaxy was found to quench at $z = 6.5$, about 800 million after the Big Bang. The phase of starburst further supports the idea of a rapid depletion of molecular gas which is needed for star formation.

However, recent discoveries have found that molecular gas can persist, remain, or even rebuild after a galaxy has quenched. This paper dives into the quenching mechanism surrounding the behavior of the molecular gas and how it is not able to collapse into stars later on.

Presence of Gas in Quiescent Galaxies

Historically, galaxy quenching has been considered a binary process. Galaxies were thought to undergo a complete shutdown process that left them unable to form stars. This theory required that all the molecular gas in the galaxy must be exhausted or expelled into the universe. The primary causes of the removal of gas were considered either AGN feedback to expel the

cold material or violent starburst to consume all of it. Consequently, quiescent galaxies were expected to be completely "dry" and contain little to no star-forming gas.

However, more recent studies demonstrate how gas depletion does not always accompany the rapid shutdown of star formation in a galaxy. A study by French et al. focused on a population of post-starburst galaxies (PSB). These galaxies represent the rapid transition state between a star-forming galaxy and a quiescent one. These galaxies are found through Balmer absorption lines, which are found through hydrogen absorbing light in medium sized stars. The Balmer absorption features reveal a burst of A-types stars (French et al.). These stars have a lifetime of around one billion years, which makes them useful as "clocks" to determine if a starburst occurred recently. Additionally, astronomers detect weak nebular omission in these galaxies, which are caused by hot gas surrounding massive, brief stars. The signature of these two confirms that star formation has rapidly shut down (French et al.).

PSBs serve as ideal galaxies to test and study how galaxies undergo structural changes and how star formation is suppressed. French et al. found that most of these galaxies retain a large amount of the interstellar medium (ISM) rather than completely being empty or consumed during the period of starburst. This example points to how galaxies can retain molecular gas even after quenching (French et al.). Additionally, it proves that the mechanisms for shutting down star formation must be significantly faster than the time required to exhaust or eject the gas reservoirs.

Additionally, mapping from the EDGE-CALIFA survey found that cold molecular gas can remain in galaxies undergoing quenching in the local universe (Colombo et al.). Through observations of carbon monoxide from the survey and observations of 472 CALIFA galaxies, Colombo et al. were able to analyze molecular gas presence in these galaxies. The cold gas was directly detected across local galaxies, which included the galaxies outside of the Star Formation Main Sequence and quenched galaxies.

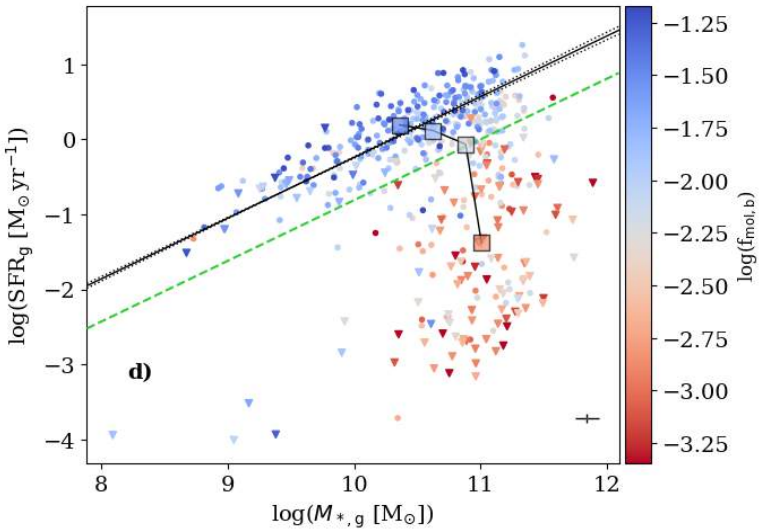


Fig. 1. The fraction of molecular gas with respect to stellar mass for the CALIFA galaxy sample, showing the spread in gas reservoirs in both active and quiescent galaxies. Adopted from Colombo et al. (their Fig. 3d)

Galaxies do not need to be completely clear of molecular gas to quench. As the galaxies begin to quench, the star formation efficiency (SFE) decreases significantly over time, meaning that the remaining molecular gas becomes much more inefficient at making stars (Colombo et al.).

Direct observations of seventeen old, massive quiescent galaxies further prove that molecular gas survives in quenched environments (Lorenzon et al.). The galaxies were 3-10 Gyr old or z ~ 0.4 and their solar masses ranged from $10^{10.4}$ and $10^{11.25}$ solar masses. Of the 17 galaxies, 10 show direct > 3σ detections, proving molecular gas was regularly able to survive quenching. These galaxies retained a molecular gas fraction ranging from 1.7% and 14%, proving that some "dead" galaxies are not completely dry of molecular gas (Lorenzon et al.).

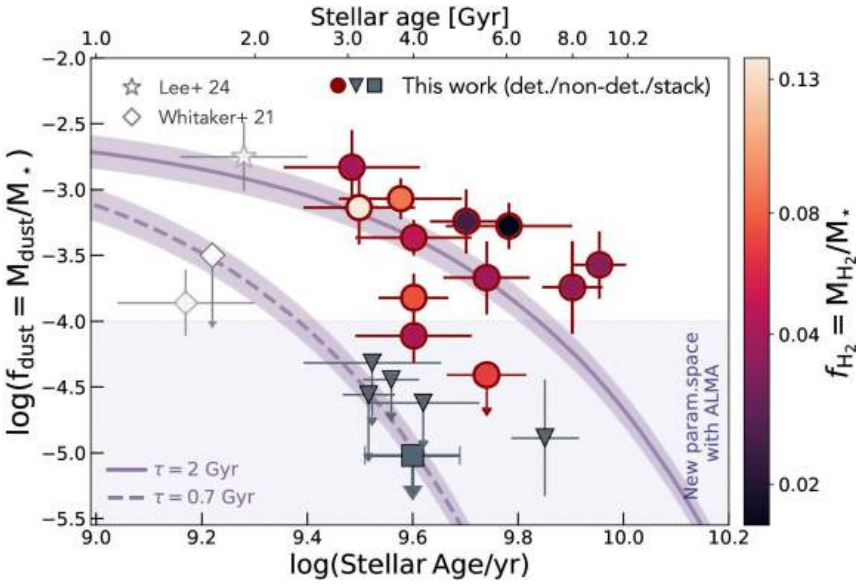


Fig. 2. Graph comparing stellar age to the dust fraction, the ratio of dust to total star mass in the galaxy as appearing in Lorenzon et al. (their Fig. 2.). The color bar represents how much cold molecular gas the galaxy has, where red represents more gas and purple represents less. The dashed and solid lines reflect simulation predictions.

In Figure 2, the points do not follow a linear pattern as some older galaxies have low dust and gas levels while others are found to contain a large amount of dust, gas, or a combination of the two. This observation proves the decoupling of molecular gas and dust after quiescence since the galaxies follow different pathways.

Furthermore, the article found that dust and molecular gas begin to evolve independently after star formation shuts down. Half of the sample shows a rapid decline of dust while the other half holds a high dust content (Lorenzon et al.). Because of this, the detection of dust cannot be used to correlate to the amount of molecular gas in quiescent galaxies, since cold molecular gas exists across diverse ISM states (Lorenzon et al.).

Another study done by Valentino et al. focused on the massive, quiescent galaxy RUBIES-UDS-QG-z7, which was observed at a redshift of z = 7.27, about 700 million years

after the Big Bang. ALMA Band 6 observations directly detected atomic carbon emission of the C II line. The galaxy keeps a cold molecular gas fraction of $> 20\%$ (Valentino et al.).

A further mapping from JWST/NIRCam found that the distribution of stars was very compact. Inversely, the C II mapping found that the cold gas is spread out, as about 70% of the total molecular gas lives outside of the body of the main galaxy, residing in a circumgalactic halo (Valentino et al.).

An analysis of JWST/MIRI mid-infrared data across about 5,000 galaxies found interstellar dust as a persistent quality of quiescent galaxies (Lisiecki et al.). About 13% of all quiescent galaxy candidates showed significant dust attenuation, meaning most quiescent galaxies still maintained a dust rich environment. Additionally, the retention of dust was long-term across the ISM, meaning that dust retention is a stable feature that does not stop shortly after quenching (Lisiecki et al.).

While Lorenzon et al. demonstrates how dust presence cannot correlate directly to cold gas presence, the interstellar dust suggests that galaxies retain an active ISM long after they have quenched rather than complete exhaustion of all ISM material which was previously accepted.

Internal and Morphological Quenching

Molecular gas remains present in some molecular galaxies, but still cannot form stars. The way gas forms stars is by collapsing under its own gravity and heats up. However, star formation is quenched even in the presence of molecular gas through morphological quenching (MQ). MQ is the process where the growth of a central stellar spheroid, a bulge, stabilizes the gas disk against gravitational collapse, stopping star formation without removing or consuming the gas (Martig et al.). The central bulge, often built via galaxy mergers or internal evolution, increases the epicyclic frequency and tidal shear in the galaxy, preventing cold gas from forming a dense cloud to collapse into a star.

MQ further explains how a “red and dead” galaxy does not require removing all cold gas in a galaxy (Martig et al.). Unlike AGN feedback and ram-pressure stripping, morphological quenching operates in lower mass halos such as $< 10^{12}$ solar masses.

Further observational data was collected on massive Green Valley (GV) galaxies at a redshift of $0.2 < z < 1.2$ (Su et al.). GV galaxies are a transitional population of galaxies located in between the blue cloud and red sequence. Blue cloud galaxies are actively star-forming while the red sequence contains “dead,” passive galaxies. Based on the COSMOS 2020 catalog, galaxies were split into early-type disks (ETD), which were dominated by a bulge and with a disk, and late-type disks (LTD), which were only disk-dominated. ETDs had shown lower sSFR than LTDs even though both groups of galaxies had very similar environments. To quantify this, Su et al. established a morphological quenching efficiency parameter (Q_{mor}). Q_{mor} was much more sensitive to stellar mass than redshift, especially with galaxies with solar mass $> 10^{10.8}$ at $z < 1$. These observations further confirm that the mass concentration of galaxies suppress star formation independently of environmental factors (Su et al.).

While bulges alter mass concentrations in a galaxy, other structures such as stellar bars, suppress star formation by altering gas kinematics (Géron et al.). Stellar bars are linear structures made of stars and gas that span across the entire galaxy (See Fig. 3). As bars rotate, it exerts a gravitational torque that drives gas either into the core or the edges of the bar.

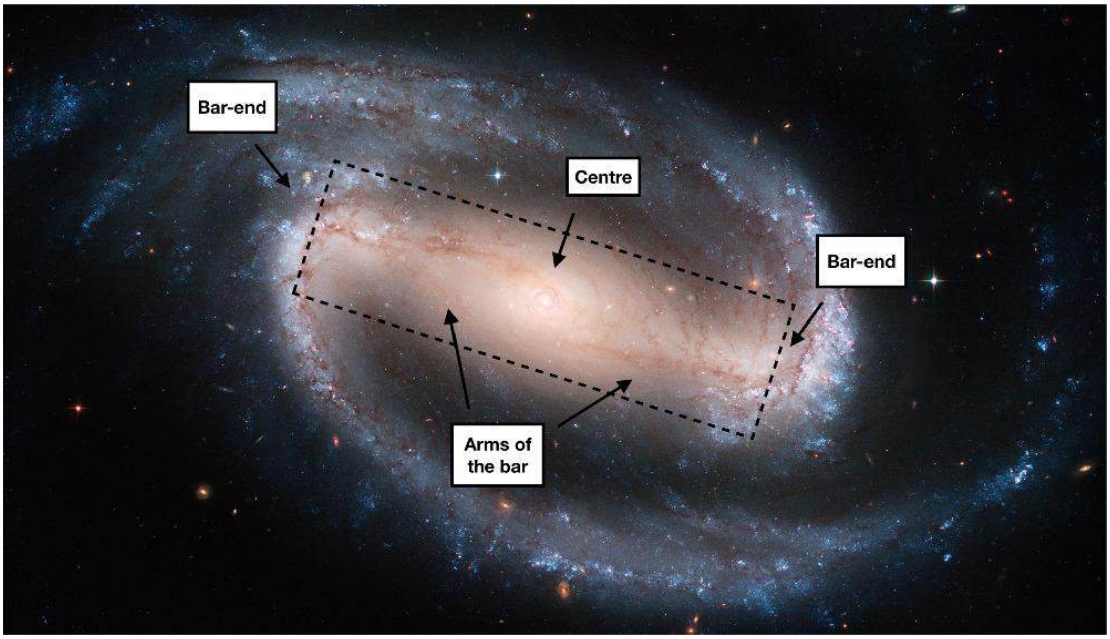


Fig. 3. A representation of stellar bars in a galaxy as appearing in Géron et al. (their Fig. 1). Image of galaxy NGC 1300, where the center of the stellar bar is concentric with the center of the galaxy. The dust bars are easily visible and the spiral connects to the ends. Credits: NASA, ESA, and The Hubble Heritage Team (STScI/AURA).

Bars were classified as strong or weak using Galaxy Zoo DESI (Géron et al.). Strong bars in active galaxies actively suppressed star formation within the arms while also temporarily enhancing star formation in the core and ends. However, weaker bars underwent star formation similar to galaxies with no bar, and total SFR did not significantly differ between the fast and slow bars, only the regions where star formation occurs changed (Géron et al.). Thus, even though stellar bars only relocate star formation, the inflow of gas to the center accelerates the growth of a central bulge to drive morphological quenching.

Further observational studies reinforce how morphological transformation correlates to star formation quenching using the SDSS DR7 catalog (Liu et al.). Galaxies were classified into four different categories: star-forming early-types (sEs), quenched early-types (qEs), star-forming late-types (sLs), and quenched late-types (qLs). Early-type galaxies featured a central bulge while late-types were disk-dominated. There are two primary pathways for quenching. First, the galaxies start as a sLs before converting into a sEs and quenching as an qEs, meaning they built a central bulge first before completely quenching. The other pathway is starting as a sLs and quenching as a qLs before converting into a qEs, meaning the galaxy quenched before a central bulge formed (Liu et al.). This study further confirmed that

morphological quenching depends highly on stellar mass and central mass density, demonstrating how structural transformations are primary drivers for quiescence.

Examination of galaxies at a higher redshift also demonstrates key morphological changes of quenched galaxies. Gozaliasl et al. utilized COSMOS-Web NIRCam imaging to map the brightest group galaxies across $0.08 < z < 3.7$. Across all redshifts, quiescent galaxies were found to be more compact than star-forming galaxies at certain stellar masses. Additionally, quiescent central galaxies experience stronger size growth from $z = 3.7$ to the local universe, reinforcing the concept that quenching pathways produce a central bulge through internal mechanisms (Gozaliasl et al.).

ISM Conditions and Feedback

Further studies investigated the structure of the interstellar medium and how the cold, molecular gas had been disturbed to prevent further star formation. Smercina et al. utilized the ALMA CO line across six PSB galaxies that were rich in both gas and dust. They found that the ISM was concentrated in extremely compact reservoirs on sub-kiloparsec levels, meaning closer than 3,260 light years. These reservoirs rival the ones present in luminous, infrared galaxies, but star formation is still highly suppressed (Smercina et al.).

The internal turbulent and kinetic pressure in the molecular gas of the PSBs was far greater than in a normal, star-forming disk. This pressure causes significant turbulent heating, disturbing the cold gas and preventing it from collapsing (Smercina et al.).

Zanella et al. observed CO (5-4) emission in eight gas rich post-starburst galaxies across the redshifts of $0.6 < z < 1.3$ to further investigate the state of molecular gas in quiescent galaxies. Galaxies previously had low-J CO detections, meaning a lower temperature while requiring little thermal energy to excite the gas. J represents rotational angular momentum quantum number, and as molecules transition from higher to lower states, they emit photons at a specific frequency (Zanella et al.). The low-J detections trace the total cold molecular gas reservoir throughout the galaxy.

However, this study was the first to detect high-J excitation in PSB galaxies in interacting systems beyond the local universe. High-J transitions require higher temperatures to excite and much higher volume densities. The high-J excitation detected traced the warm and excited gas in the quenched systems.

Moreover, isolated PSBs have a much lower line luminosity excitation ratio, meaning the excitation levels are twice lower than local star-forming disks (Zanella et al.). The low excitation ratio proves that the surviving molecular gas is not dense enough to form stars.

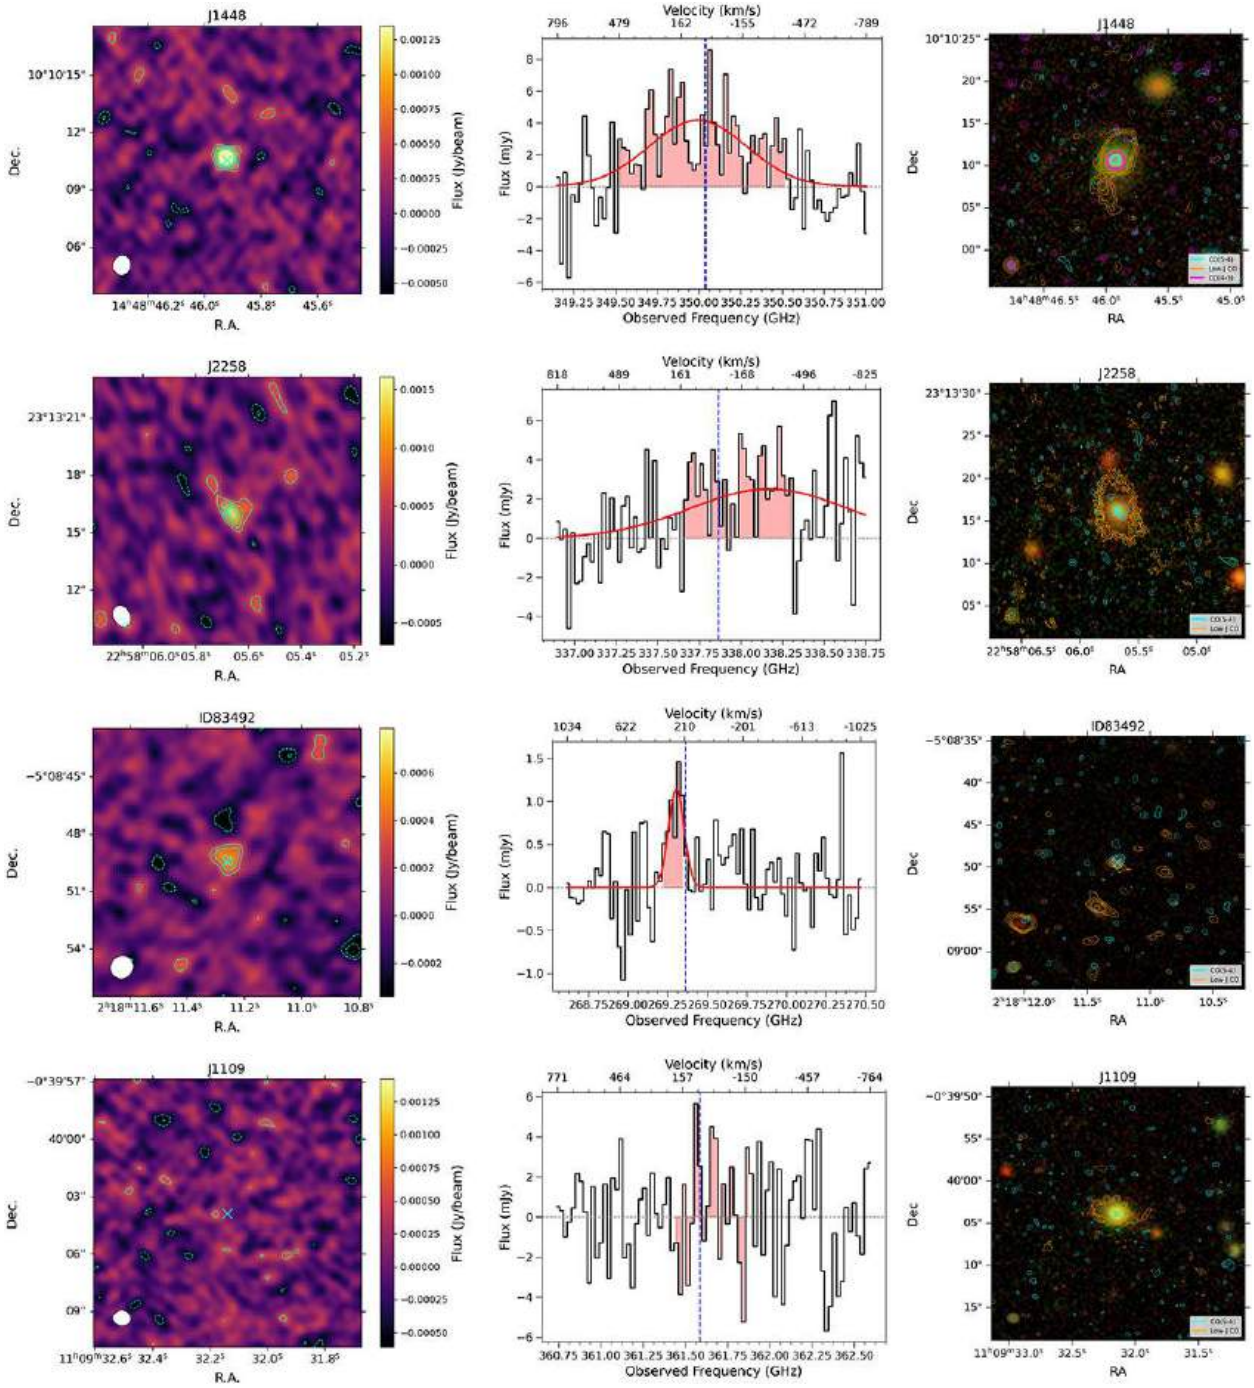


Fig. 4. ALMA data of PSBs adopted from Zanella et al. (their Fig. 1). The first column represents the CO(5-4) line and the second column displays the exact spectral line emitted by the gas. The final column represents the starlight from the galaxies on top of the CO emission.

To test if the lack of X-ray or AGN signatures in PSBs are only due to dust covering the readings, hiding a central quasar, Hewitt et al. analyzed a sample of 65 PSB galaxies at the redshift of $1 < z < 2$. Using JWST/MIRI imaging, they found that massive PSBs with a solar mass $> 10^{10}$ displayed no mid-infrared excess compared to older quiescent galaxies. The results

confirmed that post-starburst galaxies do not contain any hidden AGN activity during their transition from starburst to post-starburst (Hewitt et al.).

Analysis of cold gas excitation further confirms that galaxies remain quiescent because of a lack of dense molecular gas rather than a complete deficit of cold gas (French et al.). ALMA Band 7 observations targeted the carbon monoxide line in three PSB galaxies, which were used to compare line ratios. The galaxies were found to have low CO excitation which was similar to early-type galaxies (French et al.). The low excitation points out that lower gas density rather than high temperature suppressed star formation in the post starburst phase.

This conclusion is further confirmed by the low HCO^+/CO luminosity ratios displayed by the PSB galaxies. Further data from the SDSS MaNGA show that the gas is both spatially and kinematically disturbed. The molecular gas cannot settle into dense cores that can condense into stars as it cannot collapse under its own gravity. Instead, the gas experiences multiphase gas outflows, meaning that a combination of cold molecular gas and warm ionized gas are pushed around but not enough to clear the entire ISM (French et al.).

A previous study by Smercina et al. found that the infrared spectra of post-starburst galaxies exhibited a unique signature. They displayed an unusually strong rotational hydrogen (H_2) emission. The presence of this emission with suppressed star formation rates corroborates that the ISM is held back against gravitational collapse to form stars by turbulent heating into the gas (Smercina et al.). Turbulent heating does not mean the gas was raised to a high temperature, but was constantly "stirred" so that it is unable to collapse into a star.

Additionally, the polycyclic aromatic hydrocarbon to total infrared (PAH/TIR) ratio and the dust-to-burst stellar mass ratio decrease with age after starburst (Smercina et al.). This trend matches the effects of gas destruction as galaxies move from a gas-rich merger to a hot ISM.

Ghaffari et al. analyzed 25 rich galaxy groups from the redshift of $0.18 < z < 3.65$ in the COSMOS-Web to further understand the relationship between galaxy quenching and structure. In massive galaxies, star formation shutdown was seen to happen before morphological transformations in the galaxy. AGN feedback plays a partial role in helping cease star formation in massive central galaxies, since it injects energy into the intra-group medium (IGrM). This prevents the gas from cooling down to form stars as gravitational changes begin to occur. The morphological transformation in galaxies maintains suppressed star formation by disturbing the ISM after it has been quenched (Ghaffari et al.).

Environmental Factors and Rejuvenation

Investigating environmental drivers of quenching, Oxland et al. observed satellite galaxies across groups and clusters to track sSFR. For both group and cluster environments, star formation quenching occurs faster than morphological transformations. This concludes that environmental processes such as ram pressure stripping or starvation quench star formation on short time-scales of less than 1 Gyr (Oxland et al.). Then, structural transformation such as central bulge growth drives quenching over longer timescales.

Mapping properties in SDSS DR18, Nandi & Pandey classified environments into sheets, filaments, and clusters and constructed stellar-mass-matched subsamples of the environments. This allowed isolation between cosmic web effects and mass-dependent ones. Quenched fractions increased along the gradient, as they were lowest in sheets and highest in clusters (Nandi & Pandey). However, the fraction plateaus at solar mass of $10^{10.6}$, meaning there is a transition from environmental-driven quenching in lower masses and internal quenching at a higher mass.

However, unlike the quenched fraction, the bulge fraction rises with stellar mass across all three environments (Nandi & Pandey). This corroborates the conclusions from Oxland as it decouples star formation quenching with morphological transformations of galaxies.

At $z > 3$, dusty star-forming galaxies often have the same colors of massive quiescent galaxies, and low-mass quiescent galaxies at $z > 5$ are very hard to detect. Utilizing JWST/MIRI mid-IR imaging, Yang et al. investigated near-infrared wavelengths to constrain dust attenuation. Massive quiescent galaxies were found to experience a rapid decline as redshift increases, while low-mass quiescent galaxies have a constant number density across the redshift ranging from $z = 4\text{-}8$ (Yang et al.).

The flat number across redshifts proves that low-mass quiescent systems are not fully quenched. Instead, they experience temporary quenching episodes, meaning they have a brief shutdown of star formation which is followed by star formation again (Yang et al.). The galaxy is susceptible to rejuvenation as gas accretion continues in the future (See Fig. 5).

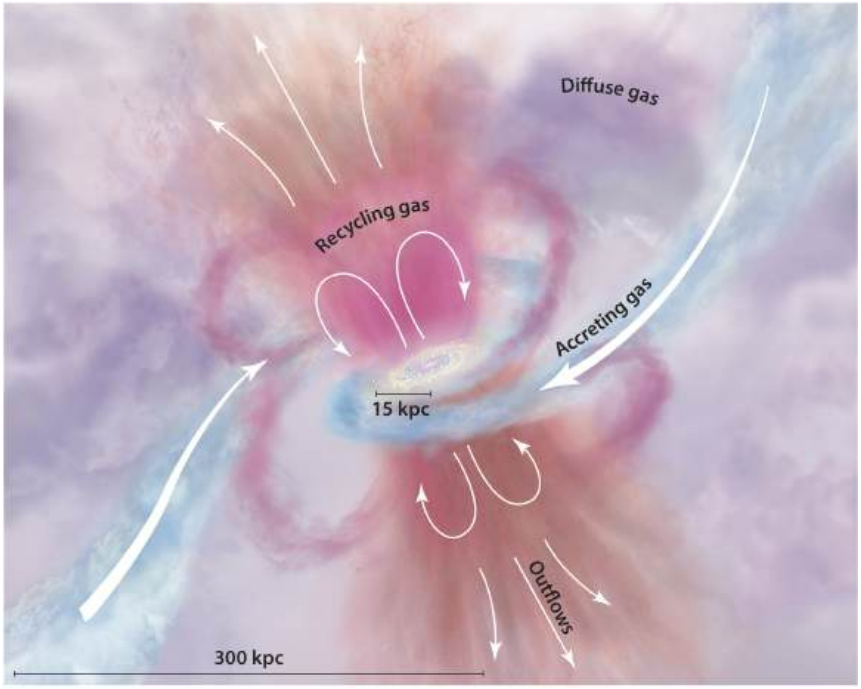


Fig. 5. Cartoon of rejuvenation as appearing in Tumlinson et al. (their Fig. 1). The galaxy's central disk is 15 kpc across. Accreting gas models cold, fresh gas flowing into the galaxy from the environment, while outflows exhaust old, unusable gas.

Beforehand, astronomers would look for a strong Hα alongside weak near-ultraviolet light (NUV). However, massive O-type stars that produce Hα also produce large amounts of NUV. To address these limitations, Lazarus & Parker used Hα readings to subtract O-star NUV emissions from total, dust-corrected NUV emissions. Across the Sloan Digital Sky Survey, about 10,000 galaxies were found to be rejuvenating (Lazarus & Parker).

Thus, about 4.5% of the local galaxy population at intermediate stellar masses were revealed. These galaxies were almost always in low-density environments and had low metallicities. This proved that rather than recycling their internal gas, the galaxies would pull fresh gas from their environment to power star formation again (Lazarus & Parker).

Direct observational evidence of star rejuvenation by Yu et al. analyzes the J1611 +4026 triple galaxy system, which consists of a massive quiescent early-type and two gas-rich, star-forming satellites. The two satellite galaxies were both undergoing tidal stripping and gas mixing during pre-processing (Yu et al.). Additionally, even though the massive host galaxy appears quiescent, they revealed a UV excess and recent starburst within the last 100 million years. The period of starburst followed a quiescent phase that lasted 6 Gyr, confirming that interactions between the satellite galaxies can transfer metal-rich gas to restart low-level star formation.

Discussion and Limitations

This study reveals that quenching mechanisms are not a singular or independent process, and some galaxies continue to change and evolve after star formation has been suppressed. For galaxies with higher masses, internal changes such as AGN feedback add energy to the ISM, heating gas and shutting down star formation as the first step of quenching. Lower mass galaxies experience factors such as ram pressure stripping or starvation, which cut off cold gas reservoirs on shorter timescales, meaning a few Gyr. While these factors do stop star formation, galaxies are able to retain the gas reservoirs and star-forming material.

Then, over time, galaxies begin the growth of a massive central bulge either through a galactic merger or the transfer of material through a stellar bar. This morphological transformation is a crucial part of stabilizing star formation for long periods of time throughout the galaxy. The central bulge stabilizes the gas and increases the gravitational potential which can work to continue low-level AGN feedback. This process of morphological quenching prevents the cold gas from collapsing into new stars as a more permanent mechanism for quiescence.

However, some lower-mass galaxies do not grow a central bulge. Through environmental factors such as tidal stripping or gas mixing, these galaxies are able to replenish their “fresh” gas reservoirs, providing material to restart star-formation. The low-level star formation occurs after a period of quiescence, allowing for rejuvenation.

While our understanding of quiescent galaxies and their interstellar medium is actively expanding, there are many data limitations currently that do not allow us to fully understand these galaxies. First, data falls short for high-redshift galaxies. As observations approach higher

redshifts, they bias toward more massive and luminous galaxies, while small or low-mass quiescent galaxies are extremely hard to detect. This makes it harder to detect rejuvenation or if quenching happens in periods rather than being permanent in low mass galaxies at high redshifts. Consequently, rejuvenation is not applicable as universal and only to some local environments. Additionally, high dust attenuation at high redshift can conceal star formation, causing active, dusty galaxies to be classified as quiescent, misrepresenting the census of quenched galaxies.

Additional observational constraints when considering the limits of ALMA and JWST. While both are necessary to map dust and investigate molecular gas tracers such as CO, ALMA noise increases greatly when pushed to higher frequencies. Detecting faint cold gas becomes much more difficult. Beyond the noise, ALMA are limited by field-of-view constraints, which makes it harder to see if a galaxy has no cold gas supply or holds a fraction of its reservoirs.

Simulations have constraints when considering spatial resolution. To properly reflect real-world cases, the physics involved with these simulations would require much more advanced black hole feedback. Currently, models struggle to simulate gas-mixing without higher numerical resolution. It is still very difficult to understand the thresholds to maintain quiescence or to allow for later rejuvenation.

Conclusion

This study explains how galaxy quenching is not a set and instant process, but a multi-stage evolution that differs based on mass, environment, and structural transformations. The point of this study was to understand the presence of molecular gas in quiescent galaxies and how the material is inhibited from collapsing into stars.

To understand this observation, the process of quenching has to be viewed less as a binary process. Quenching was thought to be considered a stage after a period of massive starburst or AGN feedback, where all molecular gas was either consumed or ejected from the galaxy. However, further observations revealed the presence of gas in massive quiescent galaxies, providing a more undiscovered part of quiescent galaxies.

Moreover, the way the molecular gas is suppressed from forming stars relates to how it remains in the galaxy after quenching. While initial star formation was shut down from AGN feedback or starvation, many galaxies also experienced the growth of a central bulge, drastically altering the gravitational potential of the interstellar medium. While the formation of stellar bars causes an initial starburst, the movement of mass toward the center and edges of the galaxy drives them to quiescence. This transformation triggered the process of morphological quenching, suppressing star formation through a change in shape. The transformation dynamically stabilized the gas discs, allowing it to remain in the galaxy without forming stars.

Further observation of the interstellar medium of the galaxies reveal the distribution of gas which prevents star formation. The gas mostly resides in outer discs rather than dense, central regions, making it much harder to form stars. This means the molecular gas is too far apart and not dense enough to satisfy gravitational collapse into a star. However, in interacting

systems, the gas is much more excited by tidal interactions and gas-mixing, making the gas warm and unable to form stars, preventing star formation even with gas present.

However, it is necessary to observe cases where molecular gas was able to collapse into stars. Certain lower-mass galaxies did not have a central bulge yet still ceased star formation. These were temporary quenching episodes as the molecular gas in these galaxies was “dirty.” Fresh gas was supplied by the environment, allowing for low-level star formation after a long period of quiescence.

To fully understand this framework and evolution of all types of quiescent galaxies, further work must expand on current data limitations. Future research through higher-resolution simulations for processes such as black hole feedback are necessary to better understand the interacting components of these processes. Tools such as the Square Kilometer Array can be used in the future to map hydrogen and radio waves across much higher redshifts. Additionally, more detailed studies with new ALMA or JWST observations can further our understanding of quiescent galaxies.

Acknowledgements

I would like to thank Dr. Ryan Farber for his mentorship, discussions, and constructive criticism for the development of this paper.

Works Cited

Baker, William M., et al. “Exploring over 700 Massive Quiescent Galaxies at z = 2-7: Demographics and Stellar Mass Functions.” arXiv.Org, 4 June 2025, https://doi.org/10.1051/0004-6361/202555829.

Carnall, Adam C., et al. “A Massive Quiescent Galaxy at Redshift 4.658.” Nature, vol. 619, no. 7971, July 2023, pp. 716–19. www.nature.com, https://doi.org/10.1038/s41586-023-06158-6.

Colombo, D., et al. “The EDGE-CALIFA Survey: Exploring the Role of the Molecular Gas on the Galaxy Star Formation Quenching.” arXiv.Org, 17 Sept. 2020, https://doi.org/10.1051/0004-6361/202039005.

de Graaff, Anna, et al. “Efficient Formation of a Massive Quiescent Galaxy at Redshift 4.9.” Nature Astronomy, vol. 9, no. 2, Feb. 2025, pp. 280–92. www.nature.com, https://doi.org/10.1038/s41550-024-02424-3.

French, K. Decker. “Evolution Through the Post-Starburst Phase: Using Post-Starburst Galaxies as Laboratories for Understanding the Processes That Drive Galaxy Evolution.” Publications of the Astronomical Society of the Pacific, vol. 133, no. 1025, July 2021, p. 072001. Institute of Physics, https://doi.org/10.1088/1538-3873/ac0a59.

———. “The State of the Molecular Gas in Post-Starburst Galaxies.” The Astrophysical Journal, vol. 942, no. 1, Jan. 2023, p. 25. Institute of Physics, https://doi.org/10.3847/1538-4357/aca46e.

Géron, Tobias, et al. “The Effects of Bar

Strength and Kinematics on Galaxy Evolution: Slow Strong Bars Affect Their Hosts the Most.” The Astrophysical Journal, vol. 973, no. 2, Sept. 2024, p. 129. Institute of Physics, https://doi.org/10.3847/1538-4357/ad66b7.

Ghaffari, Z., et al. “Star Formation Quenching Precedes Morphological Transformation in COSMOS-WEB’s Richest Galaxy Groups.” arXiv.Org, 9 Jan. 2026, https://doi.org/10.1051/0004-6361/202558283.

Gozaliasl, Ghassem, et al. “Brightest Group Galaxies in COSMOS-Web: Evolution of the Size-Mass Relation since z = 3.7.” Astronomy and Astrophysics, vol. 703, Nov. 2025, p. A129. NASA ADS, https://doi.org/10.1051/0004-6361/202556085.

Hewitt, Guillaume, et al. “No Hidden Monsters: Probing Recently Quenched Galaxies for Obscured AGN with JWST-PRIMER MIRI and NIRCam.” Monthly Notices of the Royal Astronomical Society, vol. 550, no. 2, Aug. 2026, p. stag1180. Silverchair, https://doi.org/10.1093/mnras/stag1180.

Jespersen, Christian Kragh, et al. “Explaining Ultramassive Quiescent Galaxies at $3 < z < 5$ in the Context of Their Environments.” The Astrophysical Journal Letters, vol. 988, no. 1, July 2025, p. L19. Institute of Physics, https://doi.org/10.3847/2041-8213/adeb7c.

Lazarus, Dylan, and Laura C. Parker. “An Improved UV–Optical Diagnostic for Rejuvenating Galaxies in the Local Universe and Implications for Galaxy Evolution.” The

Astrophysical Journal, vol. 998, no. 1, Feb. 2026, p. 159. Institute of Physics, https://doi.org/10.3847/1538-4357/ae2fe8.

Lisiecki, K., et al. “Impact of Stochastic Star Formation Histories and Dust Information on Selecting Quiescent Galaxies with JWST Photometry.” Astronomy & Astrophysics, vol. 708, Apr. 2026, p. A235. www.aanda.org, https://doi.org/10.1051/0004-6361/202557118.

Liu, Chenxu, et al. “The Morphological Transformation and the Quenching of Galaxies.” The Astrophysical Journal, vol. 878, no. 1, June 2019, p. 69. Institute of Physics, https://doi.org/10.3847/1538-4357/ab1ea0.

Lorenzon, G., et al. “ALMA Reveals Diverse Dust-to-Gas Mass Ratios and Quenching Modes in Old Quiescent Galaxies.” The Astrophysical Journal Letters, vol. 995, no. 2, Dec. 2025, p. L63. Institute of Physics, https://doi.org/10.3847/2041-8213/ae226c.

Martig, Marie, et al. “MORPHOLOGICAL QUENCHING OF STAR FORMATION: MAKING EARLY-TYPE GALAXIES RED.” The Astrophysical Journal, vol. 707, no. 1, Nov. 2009, p. 250. Institute of Physics, https://doi.org/10.1088/0004-637X/707/1/250.

Nandi, Anindita, and Biswajit Pandey. “Galaxy Quenching across the Cosmic Web: Disentangling Mass and Environment with SDSS DR18.” The Astrophysical Journal, vol. 1000, no. 2, Mar. 2026, p. 251. Institute of Physics, https://doi.org/10.3847/1538-4357/ae4ee4.

Oxland, M., et al. “Satellite Quenching and Morphological Transformation of Galaxies in Groups and Clusters.” Monthly Notices of the Royal Astronomical Society, vol. 529, no. 4, Apr. 2024, pp. 3651–65. Silverchair, https://doi.org/10.1093/mnras/stae747.

Skarbinski, Maya, et al. “The JWST EXCELS Survey: Insights into the Nature of Quenching at Cosmic Noon.” The Astrophysical Journal, vol. 1000, no. 2, Mar. 2026, p. 191. Institute of Physics, https://doi.org/10.3847/1538-4357/ae459d.

Smercina, A., et al. “After the Fall: The Dust and Gas in E+A Post-Starburst Galaxies.” The Astrophysical Journal, vol. 855, no. 1, Mar. 2018, p. 51. Institute of Physics, https://doi.org/10.3847/1538-4357/aaafcd.

Smercina, Adam, et al. “After The Fall: Resolving the Molecular Gas in Post-Starburst Galaxies.” The Astrophysical Journal, vol. 929, no. 2, Apr. 2022, p. 154. Institute of Physics, https://doi.org/10.3847/1538-4357/ac5d5f.

Su, Yuguo, et al. “The Impact of Morphological Quenching Mechanisms on Star Formation Activity at $0.2 < z < 1.2$ in the COSMOS Field.” Astronomy & Astrophysics, vol. 699, July 2025, p. A184. www.aanda.org, https://doi.org/10.1051/0004-6361/202553693.

Tumlinson, Jason, et al. “The Circumgalactic Medium.” arXiv.Org, 26 Sept. 2017, https://doi.org/10.1146/annurev-astro-091916-055240.

Valentino, F., et al. “Extended [CII] Gas Emission in and around a Massive Quiescent Galaxy at Z=7.3.” arXiv.Org, 19 June 2026, https://arxiv.org/abs/2606.21361v1.

Yang, Tiancheng, et al. “A Census of Quiescent Galaxies across $0.5 < z < 8$ with JWST/MIRI: Mass-Dependent Number Density Evolution of Quiescent Galaxies in the Early

Universe." The Astrophysical Journal Letters, vol. 1000, no. 2, Mar. 2026, p. L42. Institute of Physics, https://doi.org/10.3847/2041-8213/ae4803.
Yu, Yaosong, et al. "Group Pre-Processing in J1611+4026: Minor Rejuvenation of a Massive ETG Fueled by Interaction-Driven Gas Transfer." arXiv.Org, 26 Feb. 2026, https://arxiv.org/abs/2602.22904v1.
Zanella, A., et al. "Gas Excitation of Post-SB Galaxies at 0.6 < z < 1.3." Astronomy & Astrophysics, vol. 709, May 2026, p. A14. www.aanda.org, https://doi.org/10.1051/0004-6361/202558073.

Beginner's Guide to the Hubble Tension: The Basics By Kento Noda

Introduction

Even in this very moment, the Universe continues to expand outwards. Despite decades of increasingly precise observations and advances both mathematical and technological, scientists still cannot agree on a single value for the Hubble constant. The Hubble constant is defined as the rate at which the Universe expands. Although it is called a "constant," different methods for calculating this value currently yield different results. On one side is the Planck satellite, which measured the CMB with extraordinary precision and, using the ΛCDM model, derived a Hubble constant of about 67.4 km/s/Mpc. On the other hand is the SH0ES collaboration, which measured the expansion rate directly using stars and supernovae, arriving at about 73.0 km/s/Mpc. This reveals a central problem at the heart of cosmology. This paper explains the origin of this so-called "Hubble tension" by examining how the Hubble constant is measured using both early-Universe factors, such as the Cosmic Microwave Background (CMB), and late-Universe methods. The paper also describes the details of each method, starting with what the CMB is and how it is applied to calculating the Hubble constant.

CMB

The CMB is the electromagnetic radiation released approximately 380,000 years after the Big Bang. Before this time, the Universe was a hot, dense plasma, which meant that energetic photons could not travel freely. Rather, they were constantly scattered by free electrons. As the Universe expanded and cooled, electrons and protons combined to form neutral hydrogen atoms through a process called recombination, which allowed light to travel freely for the first time. The CMB is a snapshot of the first released light, and it serves as a way for people today to observe the Universe at that early period.

Today, the CMB is observed at a temperature of about 2.7 Kelvin, far colder than the original temperature of roughly 3,000 Kelvin at the time of emission. This dramatic cooling was caused by cosmic expansion. As space stretches, the energy of the photons is distributed over a greater volume, which causes their wavelength to stretch and their effective temperature to drop. It is important to note that the CMB is observed in every direction across the sky because light was emitted uniformly through the early Universe in all directions at the same time.

One of the most important features of the CMB is that it is nearly uniform. Across the full sky, the temperature fluctuations exist at a size of roughly one part in 100,000. These fluctuations are not just noise. They carry important information because of their origin, during an early period of expansion, known as inflation. This event created small density variations in the plasma. Over time, gravity amplified these variations, which eventually led to the large-scale structure of the Universe in the form of galaxies and galaxy clusters we observe today. Full-sky temperature maps of the CMB, produced by experiments such as COBE, WMAP, and the Planck satellite, have charted the fluctuations with better resolution over time. Each successive mission revealed finer details of the structure of the early Universe.

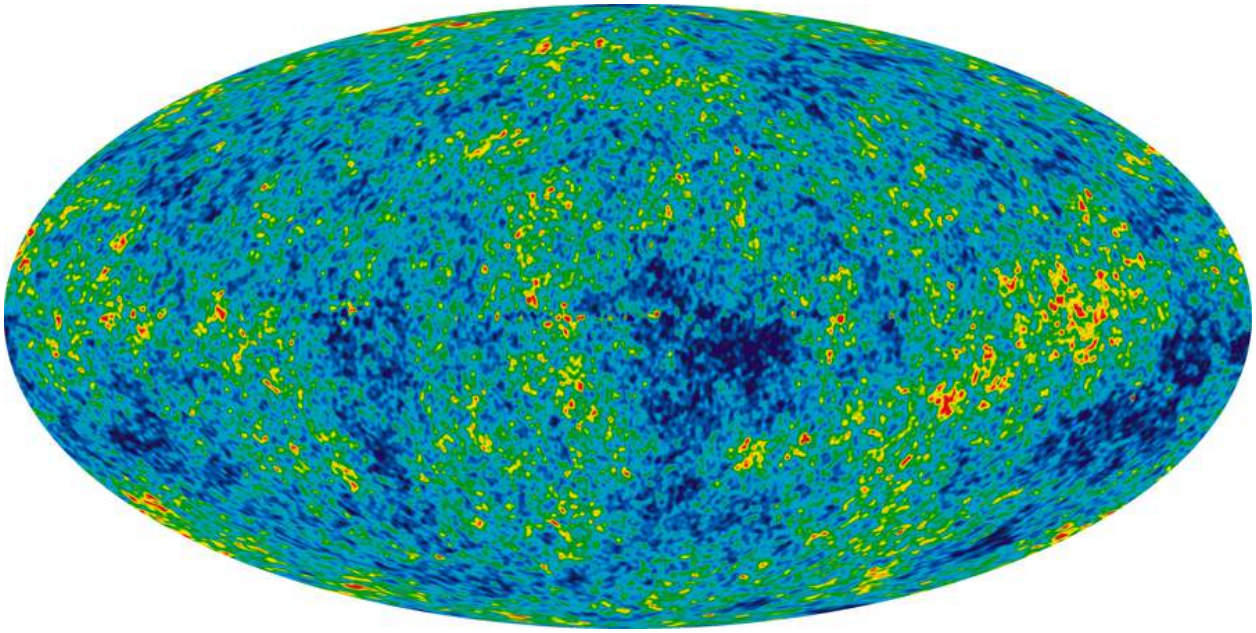


Figure 1. WMAP full-sky map of the Cosmic Microwave Background temperature fluctuations, illustrating primordial density variations.

One of the most important features of the CMB is that it is nearly uniform. Across the full sky, the temperature fluctuations exist at a size of roughly one part in 100,000. These fluctuations are not just noise. They carry important information because of their origin, during an early period of expansion, known as inflation. This event created small density variations in the plasma. Over time, gravity amplified these variations, which eventually led to the large-scale structure of the Universe in the form of galaxies and galaxy clusters we observe today. Full-sky temperature maps of the CMB, produced by experiments such as COBE, WMAP, and the Planck satellite, have charted the fluctuations with better resolution over time. Each successive mission revealed finer details of the structure of the early Universe.

More recently, ground-based telescopes have provided maps in even greater detail, though in a different way. The Atacama Cosmology Telescope (ACT), located in the Chilean desert, and the South Pole Telescope (SPT), located in Antarctica, do not map the full sky the way Planck did. Instead, they focus on smaller areas of the sky and resolve the CMB on finer scales. These experiments have added a new layer of precision on top of what Planck had already established, and their measurements have become an important part of how cosmologists refine the ΛCDM model.

Power Spectrum

Before recombination, the Universe existed as a photon-baryon fluid. In this form, gravity compressed the fluid inward, while radiation pressure pushed outward, causing acoustic oscillations, or in other words, sound waves that spread through plasma. When recombination

occurred, these waves froze in place, creating patterns in the CMB as acoustic peaks. These peaks appear at an angular scale that corresponds to oscillations that were either at maximum compression or expansion at the time of freezing. The first and most prominent peak, at an angular scale of about one degree ($\ell \approx 200$), represents the largest wave that had enough time to complete one single compression before recombination. "ℓ" is the multipole moment, a number used to describe angular scales on the CMB sky. The relationship is roughly that an angular scale of θ degrees corresponds to $\theta \approx 180°/\ell$. So $\ell \approx 200$ corresponds to an angular scale of about one degree. Instead of measuring the angles in the sky directly, cosmologists break up the CMB temperature into harmonics. This is essentially a way to break the sky's pattern into waves of varying sizes. A higher ℓ corresponds to smaller angular scales, while a lower ℓ corresponds to larger ones. The second peak is the harmonic of the first, and its existence confirms that the medium behaves like an oscillating fluid. Beyond their positions, the peaks also contain physical information. The location of the first peak tells us about the curvature of the Universe. The heights of the peaks relative to each other reveal the baryon density, since more baryons increase the amplitude of the odd-numbered peaks. All of these features relate to the sound horizon. The sound horizon is the maximum distance acoustic waves could have traveled before recombination. This sets the fundamental scale, which is used to measure peaks.

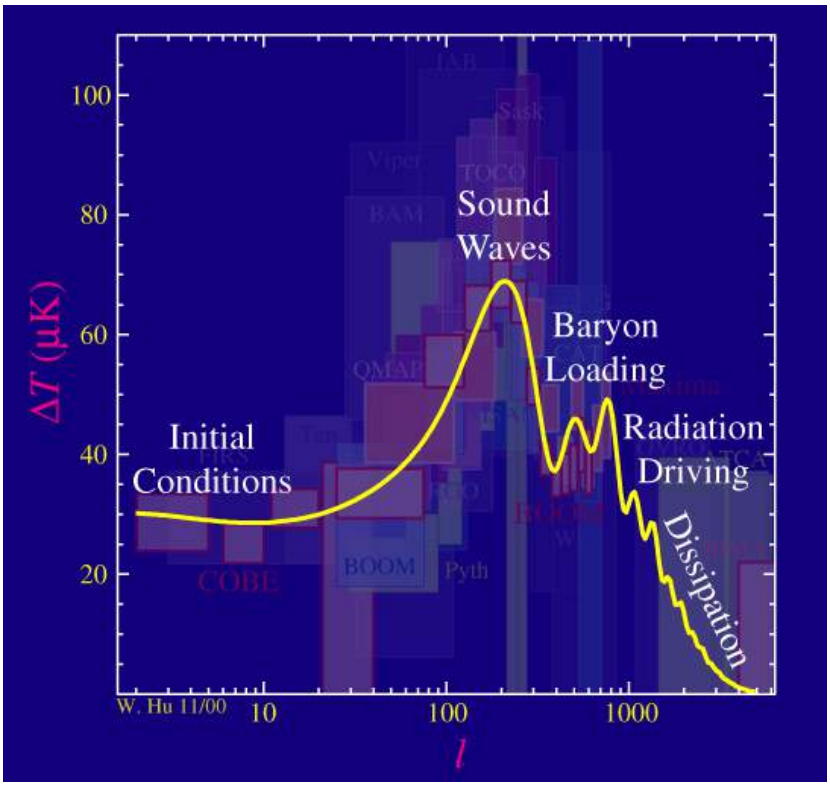


Figure 2. The CMB angular power spectrum showcasing the primary acoustic peak at $\ell \approx 200$ and subsequent sub-components.

Cosmological Parameters and the ΛCDM Model

The ΛCDM model is the standard framework used to describe the structure and evolution of the Universe. It is based on General Relativity and takes into account three major components: ordinary (baryonic) matter, dark matter, and dark energy, which is represented by the cosmological constant Λ. They make up approximately 5, 27, and 68 percent of the Universe's total energy budget, respectively. The expansion of the Universe in the ΛCDM model is governed by the Friedmann equations. The equations relate the rate of expansion to matter density, radiation, dark energy, and the curvature of space. The model also has six parameters that have to be measured rather than calculated from first principles. These include the Hubble constant H0, the density of baryons, the density of cold dark matter, the strength and shape of the primordial fluctuations, and the optical depth to reionization. When fit to Planck satellite data, the ΛCDM model determines the Hubble constant to be approximately 67.4 km/s/Mpc. This means that for every megaparsec of distance between two points in the Universe, those points are moving away from each other at about 67.4 kilometers per second. Furthermore, the fit also suggests that the Universe is geometrically flat and about 13.8 billion years old.

Late-time Probes of Hubble Constant

There is also a completely different approach to measuring the Hubble constant that does not rely on the CMB or the ΛCDM model. Instead, it uses the cosmic distance ladder, which uses a series of calibrated steps that allow astronomers to measure distances directly. To do this, they use special stars called Cepheid variables. These stars pulse at a rate that reveals how luminous they truly are, so by comparing their true brightness to their observed brightness from Earth, astronomers can calculate how far away they are. For further observations, astronomers rely on Type Ia supernovae, which are stellar explosions that always reach the same peak brightness. Once again, by comparing how bright something actually is with how bright it appears from Earth, the distance can be calculated. When this method is applied for the calculation of the Hubble constant, it results in a value of approximately 73.0 ± 1.0 km/s/Mpc. This value is meaningfully higher than the CMB-based value from Planck under the ΛCDM model, about 67.4 ± 0.5 km/s/Mpc. Because the uncertainty ranges do not overlap, the disagreement cannot be explained by the quoted statistical error bars alone. Additional late-Universe measurements, including those using Baryon Acoustic Oscillations, broadly support this higher value.

Hubble Tension

The disagreement between the two Hubble constant values, $H0 \approx 67.4 \pm 0.5$ km/s/Mpc from the early Universe and $H0 \approx 73.0 \pm 1.0$ km/s/Mpc from the late Universe, is called the Hubble tension. The gap comes to about 9%, which may not sound large at first. However, the error bars make clear that the two values do not overlap even within several standard deviations. This makes it very unlikely that the difference came from a random measurement error, pushing the statistical significance to a 5σ threshold—meaning the likelihood of it being a fluke is only 1 in 3.5 million.

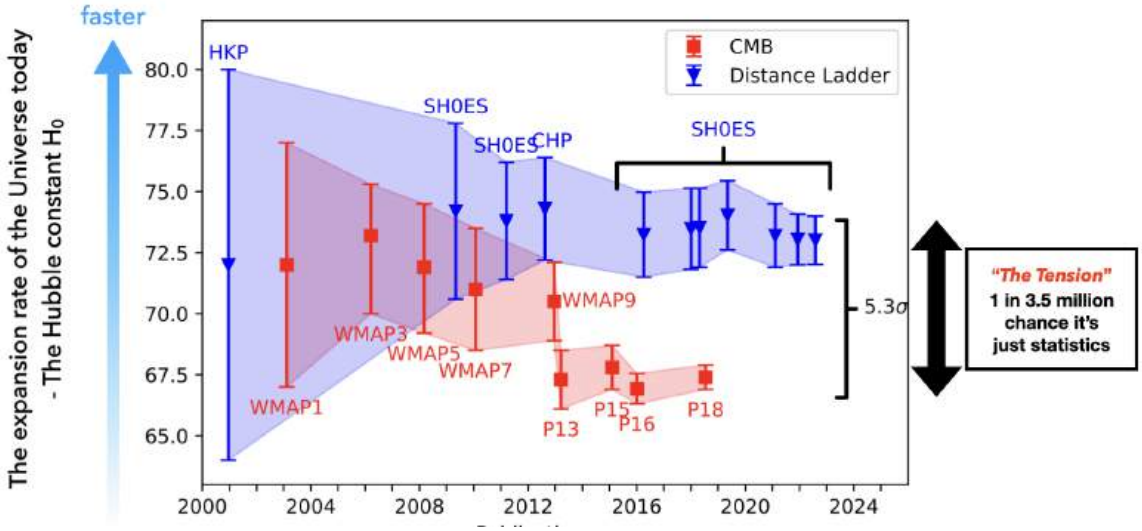


Figure 3. Historical progression of Hubble constant H_0 measurements from 2000 to 2024. The shrinking error bars of early-Universe CMB measurements (red) and late-Universe direct distance ladder probes (blue) reveal a statistically significant 5.3σ divergence known as the Hubble tension.

A great deal of research has been done to determine whether simple mistakes or biases in the measurements could explain it. For example, some teams have recalibrated the distance ladder using infrared observations, which are less affected by factors like dust. Multiple independent research teams have tried using different instruments and methods, but continue to arrive at the same high value. On the other side, independent calculations of the CMB have also had consistently low values. The fact that both sides of the disagreement hold up under repeated trials makes it increasingly difficult to say that the tension is caused by a straightforward experimental error.

There have been proposed models that can partially reduce the tension, but most of them create inconsistencies somewhere else. A possible implication is that either the ΛCDM model is missing some unknown physics, or one of our fundamental assumptions about the Universe is incorrect

Proposed Solutions to Hubble Tension

Scientists have proposed many possible explanations for the Hubble tension, and they each fall into five broad categories.

I. Early-Universe Solutions

One category proposes to change the physics of the Universe before recombination. The idea is to shrink the sound horizon, which is the distance early acoustic waves traveled before freezing. The sound horizon acts as a ruler used to figure out the Hubble constant. If something caused the early Universe to expand a little faster than expected right before recombination,

acoustic waves would have had less time to travel before they froze into the CMB pattern. That would make the sound horizon smaller, meaning the "standard ruler" used to infer cosmic distances would be shorter. With a shorter ruler, for the observed angular size of the CMB peaks to stay the same, the model would infer a larger Hubble constant, closer to what we measure today. In attempts to explain this faster expansion, scientists have proposed things like a brief burst of extra energy (called Early Dark Energy), or extra light particles that could have caused a temporary speed-up (dark radiation). The catch to this solution is that tweaking parameters of the early Universe tends to mess up many other things that our current model explains successfully. Furthermore, these fixes usually only close a part of the gap and do not come close to a full solution.

II. Late-Universe Solutions

Dark energy is the mysterious force driving the Universe's accelerating expansion. In our current ΛCDM model, this dark energy is treated as a fixed, unchanging quantity. However, some datasets hint that dark energy might be evolving. For example, Quintessence models propose that dark energy behaves more like a field that evolves, shifting as the Universe ages. Phantom energy models go further, allowing a parameter w, which describes the relationship between dark energy's pressure and density, to be low, which would indicate that dark energy density actually increases over time. These changes can alter the expansion history in ways that shift the Hubble constant, but they do not agree with some of the current data, namely, supernovae and the BAO datasets.

III. Modified Gravity

Everything we know about gravity comes from testing it in locations that are relatively nearby, usually within our Solar System or in the vicinity of merging black holes. But the Hubble tension involves the behavior of the entire Universe across billions of light-years. That is why some researchers suggest that one of our fundamental theories, General Relativity, may be inapplicable or break down on cosmological scales. This would mean that a modified theory of gravity is needed. Scalar-tensor theories and f(R) gravity models change the way gravity relates to spacetime, altering the expansion of the Universe. The challenge in this is that Einstein's theory has passed every experimental test we have thrown at it, so any modification has to be carefully designed in a way that only shows up at cosmological scales, while leaving everything else untouched, including gravitational wave observations. This is a very narrow target to hit.

IV. Dark Sector Interactions

Another class of models suggests interactions between dark matter and dark energy, allowing energy to pass between them as the Universe ages. In physics, the word "dark" usually means that something does not emit, absorb, or reflect light in a way we can directly observe with telescopes. Dark matter is inferred through its gravitational effects on galaxies and large-scale structure, while dark energy is inferred from the accelerating expansion of the

Universe. Although both names contain the word “dark,” these parameters have no known connection and are considered completely separate. If energy is slowly flowing from one to the other, the expansion history of the Universe would be different from what we currently calculate, potentially explaining the tension. The downside to these models is that they tend to have many adjustable parameters, making it difficult to constrain them with our current data.

V. Systematic Errors

Finally, it is still possible, although increasingly unlikely, that the tension is simply a result of undetected measurement errors. Maybe there is a bias in the way astronomers calibrate Cepheid stars, or an unexamined assumption in how the CMB is analyzed. Scientists take this possibility seriously and have spent years testing measurements. Yet, in the end, multiple independent teams, using different instruments, different data sources, and different methods, continue to result in the same high value for the Hubble constant. For a measurement error to be the answer, it would have to be hiding across many separate experiments and also be large, which is looking less likely the more measurements are repeated. The honest answer is that none of these solutions has been fully proven. Most of them can reduce the tension slightly, but none of them resolve it fully without creating new problems elsewhere. But then again, that is what makes the Hubble tension so exciting for researchers and cosmologists. It is a crack in an otherwise successful model of the Universe, and whatever fills this crack may teach us something new about how the cosmos works.

Conclusion

The Hubble tension refers to the discrepancy between measurements of the Universe’s expansion rate. Measurements based on the CMB and the ΛCDM model predict a lower value than observations using Cepheid variables and Type 1a supernovae. This paper also examines several proposed explanations for the Hubble tension. Because the ΛCDM model has successfully explained many cosmological observations, the persistence of the tension shows that there may be missing physics or incorrect assumptions within our current understanding. Future research will continue testing both observational methods and theoretical methods, with new experiments having more precise calibrations. This may help us determine whether the tension is caused by experimental error or hints at new physics.

Works Cited

Hu, Wayne. “An Introduction to the CMB.” Wayne Hu’s Tutorials, 1996, https://background.uchicago.edu/~whu/beginners/introduction.html.

Hu, Wayne. “Intermediate Guide to the CMB.” Wayne Hu’s Tutorials, 2001, https://background.uchicago.edu/~whu/intermediate/intermediate.html.

NASA Goddard Space Flight Center. “Mystery of the Universe’s Expansion Rate Widens with New Hubble Data.” NASA Science Missions, 25 Apr. 2019, https://science.nasa.gov/missions/hubble/mystery-of-the-universes-expansion-rate-widens-with-new-hubble-data/.

NASA Goddard. “Expansion Rate: The Hubble Tension.” YouTube, 12 Sept. 2023, https://www.youtube.com/watch?v=Jam61H0SA3Q.

NASA LAMBDA GSFC. “The History of the Universe: Graphic History of Universe Evolution.” Legacy Archive for Microwave Background Data Analysis, https://lambda.gsfc.nasa.gov/education/graphic_history/univ_evol.html.

NASA Science. “The Hubble Constant and Tension.” Hubble Science: Science Behind the Discoveries, https://science.nasa.gov/mission/hubble/science/science-behind-the-discoveries/hubble-constant-and-tension/.

National Aeronautics and Space Administration. “WMAP Overview.” NASA Science, https://science.nasa.gov/mission/wmap/wmap-overview/.

Shah, Pranav, Pedro Lemos, and Ofer Lahav. “A Comprehensive Review of the Hubble Constant Tension.” The Astronomy and Astrophysics Review, vol. 29, no. 1, 2021, article 9, https://doi.org/10.1007/s00159-021-00137-4.

Association of Universities for Research in Astronomy. “Our Mysterious Universe Still Evades Cosmological Understanding.” AURA, 6 Mar. 2023, https://www.aura-astronomy.org/blog/2023/03/06/our-mysterious-universe-still-evades-cosmological-understanding/.

Does ZEV Adoption Improve Regional Air Quality? Empirical Evidence from Los Angeles County (2010–2025) By Yinuo Jiang and Kaixi Chen

Abstract

Tailpipe emissions are major contributors to adverse respiratory health outcomes, including asthma, breathing difficulties, heart disease, and even premature death. Although the federal tax credit for new zero-emission vehicle (ZEV) purchases ended in September 2025, the California state government has recently continued to offer its own incentives for new ZEV purchases. It sparked debates over whether these programs are worth the cost and whether they actually improve air quality in the state. While past research has shown that the transition to EVs has yielded measurable improvements in air quality and health, it has mainly focused on NO2 as the primary pollutant because of the greater availability of monitoring sites and of health impacts already shown to be NO2-related. This study explored the correlation between ZEV adoption rate and changes in a broader range of pollutants, namely NO2, NO, CO, PM2.5, PM10, SO2, and lead, between 2010 and 2025. The results will contribute to answering questions like "Is Governor Newsom's EV policy worth it?" "Which pollutants have been most significantly reduced?" It will also provide a fuller understanding to inform future ZEV policy decisions.

1. Introduction

Under the California Air Resources Board (CARB) Advanced Clean Cars II (ACC II) regulation (codifying Governor's Executive Order N-79-20), all new light-duty passenger cars, pickup trucks, and SUVs sold in California must be 100% zero-emission vehicles (ZEVs) by 2035.

After the federal tax credit of up to $7,500 for new zero-emission vehicle (ZEV) purchases ends on September 30, 2025, the California state government swiftly announced it will continue to offer its own incentive policy for new ZEV purchases. In July 2026, California Governor Gavin Newsom signed Senate Bill 168 to establish the MyFirstEV program, a $270M zero-emission vehicle (ZEV) program that offers first-time buyers an instant point-of-sale rebate of $3,500 off new EVs and $1,750 off used ones. (Ohnsman 2026). It consists of $135.5 million in California state funds matched dollar-for-dollar by 13 participating automakers. (Ohnsman 2026). With a $50,000 MSRP cap, it also favors lighter passenger car models. (Roy 2026). Although these cost savings will not solve every problem standing between California drivers and electric cars, they will help, experts say (Roy 2026).

Meanwhile, California provides major point-of-sale incentives for electric commercial trucks through the California Hybrid and Zero-Emission Truck and Bus Voucher Incentive Project (HVIP). (California HVIP) It offers point-of-sale discounts on zero-emission trucks and buses, making it easier and more affordable for California fleets and businesses to switch to cleaner transportation. (California HVIP). So far, the HVIP program has provided over $1 billion through vouchers. Additionally, on May 13, 2026, Governor Newsom announced a new $1

billion California Clean Fuel Reward (CCFR) Program. It is a new rebate program for electric medium- and heavy-duty trucks. (Governor Gavin Newsom).

However, various policy experts have criticized this program, citing California's multibillion-dollar budget deficit. Claiming that spending scarce public funds on vehicle rebates is structurally unsustainable because it covers too small a share of the market to have an immediate impact. It also risks giving bonus incentives to wealthier buyers or large corporate fleets instead of funding low-income communities or struggling local transit agencies.

Set aside the argument over whether the policy is effective at encouraging more lower-income first-time buyers or small business owners to pull the trigger, the question remains whether EV adoption plays a significant role in mitigating pollution.

Numerous studies have shown that tailpipe emissions are major contributors to adverse respiratory health outcomes, including asthma exacerbations, reduced lung function, cardiovascular disease, and premature death (Garcia et al., 2023). These harmful pollutants include ground-level ozone (O3), fine particulate matter (PM2.5, particles less than 2.5 micrometers in diameter), coarse particulate matter (PM10, particulate matter that has a width or diameter of 10 micrometers or less), nitrogen dioxide (NO2), nitrogen monoxide (NO), carbon monoxide (CO), sulfur dioxide (SO2), and lead (Pb).

Heavy-duty diesel trucks are a particularly significant concern for air quality because they are the largest mobile source of PM2.5 in urban areas (Heinrich and Wichmann, 2004; Nel et al., 2001; Seagrave et al., 2006). While light-duty vehicles have benefited from increasingly stringent emission controls over the past two decades, heavy-duty trucks have longer fleet turnover rates. They continue to emit very high levels of PM2.5 and associated toxic emissions per vehicle mile traveled (Beverly, 2005). Communities adjacent to major freight corridors and ports, where dense truck traffic converges with residential neighborhoods, were significantly affected. (Beverly, 2005). PM2.5 concentrations dropped to baseline levels within 150 to 500 meters downwind of major roadways (Hitchins et al., 2000; Zhu et al., 2002), indicating that residents living in this impact zone will be exposed to higher levels of diesel-related pollutants. Past research has confirmed a correlation between PM2.5 exposure and adverse health outcomes, including asthma exacerbations, cardiovascular disease, and premature death (Delfino, 2002; Gauderman et al., 2007; McConnell et al., 2006; Sarnat and Holguin, 2007).

Past research has evaluated whether the transition to EVs has yielded measurable improvements in air quality and health benefits. Garcia et al. (2023) studied the correlation between ZEV adoption, air pollution, and health across California at the zip code level from 2013-2019 using DMV registration data, EPA NO2 monitoring data, and asthma-related ER visits. They found that although NO2 reductions were small, for every additional 20 ZEVs per 1,000 people, asthma-related ER visits dropped by 3.2%. (Garcia et al. 2023) The study also revealed a significant socioeconomic gap in EV adoption, with higher-education areas adopting EVs at substantially faster rates than lower-education areas. (Garcia et al. 2023)

Adding to this, Eckel et al. (2026) performed a longitudinal study tracking NO2 levels across California from 2019 to 2023 with TROPOMI satellite data. According to the findings,

every increase of 200 ZEV per ZCTA is correlated with a 1.10% decrease in the total annual average NO2. (Eckel et al. 2026)

Another third-line study in Taiwan used a regional chemical model to examine gaseous and particulate pollutants. They examined the impact of transitioning to EVs. They analyzed the impact of emissions released in the power grid that generates energy for EVs in the first place. (Li et al. 2016)

To cover more conditions, a study by Andre examined emissions by simulating the emission process under different conditions: renewal of diesel vehicles, promotion of petrol or electric vehicles, and traffic reduction. (Andre et.al.) They studied how these conditions could affect the amount of emissions being released.

Despite these contributions, existing research has mainly focused on NO2 as the primary pollutant because of the greater availability of monitoring sites and of health impacts already shown to be NO2-related. It leaves an open space to fill regarding how EV adoption affects other pollutants, such as CO, PM2.5, and NO. With a lack of a more detailed examination of different pollutants, our understanding of the full environmental impact of EV adoption is incomplete.

This study aims to explore the correlation between ZEV adoption rate and changes in a broader range of pollutants, namely NO2, NO, CO, PM2.5, PM10, SO2, and lead, between 2010 and 2025 in Los Angeles County, the most populous county in the US. The results will contribute to answering questions like “Is Governor Newsom’s EV policy worth it?” “Which pollutants have been most significantly reduced?” It will also provide a better understanding to inform future ZEV policy decisions.

2. Data

2.1 EV Sales Data

The dataset assesses previous ZEV indicators, including annual registration of new vehicles and the population of all on-road vehicles in Los Angeles County from 2010-2025. Light-duty vehicles are differentiated by fuel type for annual registration, including BEV (Battery Electric Vehicles), PHEV (Plug-in Hybrid Electric Vehicles), and FCEV (Fuel Cell Electric Vehicles). These data are collected from the California Energy Commission (CEC) and the California Department of Motor Vehicles (CA DMV). The CEC’s Zero-Emission Vehicle and Infrastructure Statistics were compiled in partnership with the CA DMV, which provides the registration database used to identify fuel types via VIN decoding.

Table 1. Historical ZEV Sales Record by Fuel Type in Los Angeles County (2010–2025)

Year	BEV	PHEV	FCEV	Total ZEV Population	Annual New ZEV Sales	New ZEV Sales Market Share (%)
2010	120	125	5	250	200	0.1
2011	900	1,080	20	2,000	1,500	0.4
2012	3,120	3,315	65	6,500	4,500	1.1
2013	7,500	7,350	150	15,000	9,200	2.1

2014	14,560	13,160	280	28,000	13,500	2.9
2015	23,320	20,240	440	44,000	16,400	3.4
2016	35,200	28,160	640	64,000	20,100	4.2
2017	52,200	36,900	900	90,000	27,300	5.8
2018	76,800	49,920	1,280	128,000	40,900	7.6
2019	101,680	60,680	1,640	164,000	43,200	8.1
2020	128,640	70,350	2,010	201,000	41,800	9.4
2021	176,880	88,440	2,680	268,000	72,500	14.2
2022	243,440	110,980	3,580	358,000	101,400	21.3
2023	332,500	137,750	4,750	475,000	132,600	25.8
2024	421,200	157,950	5,850	585,000	138,100	26.1
2025	507,350	180,700	6,950	695,000	142,500	27.4

Source: California Energy Commission (CEC) and the California Department of Motor Vehicles (CA DMV).

Table 1 shows the historical ZEV sales record by fuel type in Los Angeles County from 2010 to 2025. Figure 1 shows historical ZEV adoption and new vehicle sales market share from 2010 to 2025 in Los Angeles County. From Figure 1, we can roughly categorize the 16-year progress into three phases.

- Phase 1 (2010-2018): Annual sales grew from 200 units to more than 40,000 units. The total on road and ZEV population was above the 100,000 baseline. Phase 1 indicates a steady baseline growth.
- Phase 2 (2019–2023): ZEV adoption increased dramatically during this phase. The annual sales of ZEVs increased from 41,800 units in 2020 to 132,600 units in 2023, which is a 217% increase in just three years.
- Phase 3 (2024-2025): The growth rate has dramatically slowed down, and the market has stabilized. By the end of 2025, the number of on-road ZEVs in Los Angeles County is estimated at 695,000, which helps explain why additional incentives are necessary to further boost the EV adoption rate.

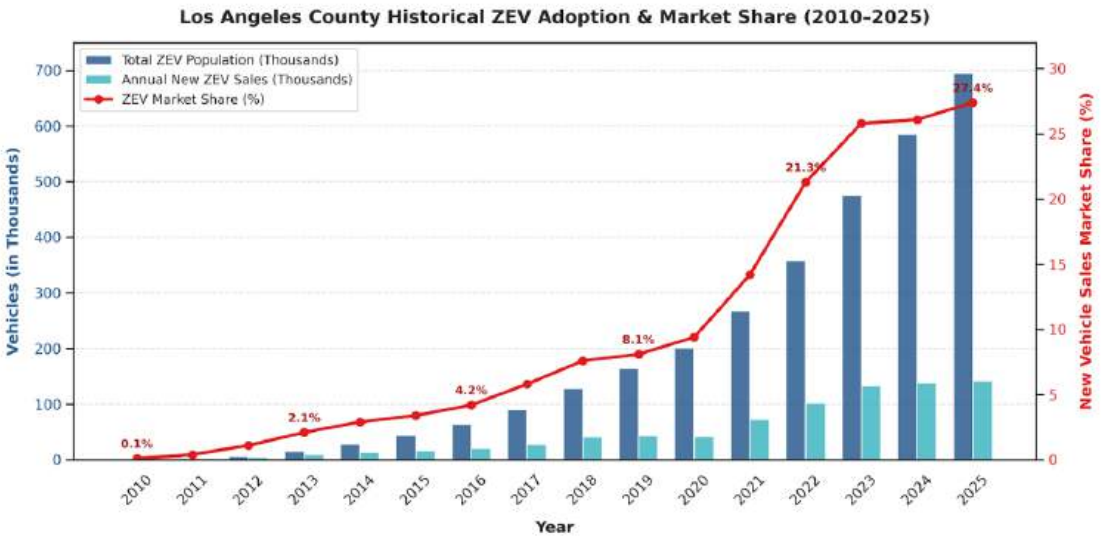


Fig 1: Los Angeles County Historical ZEV Adoption and New Vehicle Sales Market Share (2010-2025). Note: Image generated by Google, 2026 (https://gemini.google.com/).

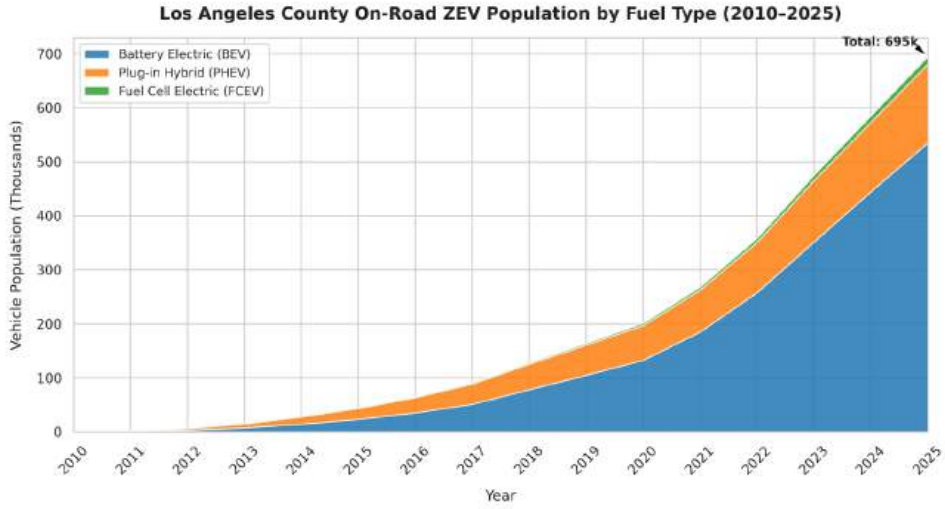


Fig 2: Los Angeles County Light-Duty ZEV Population by Fuel Type (2010-2025). Note: Image generated by Google, 2026 (https://gemini.google.com/).

Figure 2 shows Los Angeles County's light-duty ZEV population by fuel type. It highlights a significant expansion of the market for these vehicles. The shift was initially between PHEVs (Plug-in Hybrid Electric Vehicles) and BEVs (Battery Electric Vehicles). Over the last decade, this shift indicates a transition towards pure electric powertrains. By early 2026, BEVs make up about 73% of the on road ZEV population. This shift shows BEVs dominance in the market, which also reflects the utility scale of infrastructure, longer ranges, and changes in consumers' purchase preferences.

2.2 Air Pollutants Data

This dataset integrates records from the SCAQMD Historical Data Tables (2010–2025), EPA AirData Design Value Reports, and CARB's iADAM Database.

Table 2: Air Pollutants Trend in Los Angeles County (2010–2025)

Year	Ozone (O3)*	PM2.5 (µg/m3)	PM10 (µg/m3)	NO2 (ppb)	CO (ppm)	SO2 (ppb)	Lead (Pb) (µg/m3)
2010	123	14.9	26.5	22.1	3.2	12.4	0.05
2011	122	14.5	27.1	21.8	3.1	11.8	0.04
2012	138	14.1	28.2	21.2	3	11.2	0.04
2013	116	13.8	29.5	20.8	2.9	10.5	0.03
2014	126	13.5	31.4	20.4	2.8	9.2	0.03
2015	113	13.1	30.5	19.8	2.7	8.2	0.02
2016	132	12.8	29.7	19.2	2.6	7.1	0.01
2017	145	12.6	30.2	18.8	2.5	6.2	0.01
2018	141	12.5	30.1	18.2	2.4	4.6	<0.01
2019	126	12.3	29.8	17.5	2.3	3.9	<0.01
2020	157	14.4	36.8	16.8	2.2	3.5	<0.01
2021	130	13.1	30.5	16.5	2.1	3.2	<0.01
2022	124	12.4	29.1	15.8	2	3	<0.01
2023	118	11.9	28.5	15.2	1.9	2.8	<0.01
2024	121	12	29	14.8	1.8	2.6	<0.01
2025	115	11.8	28.7	14.5	1.7	2.5	<0.01

Source: SCAQMD Historical Data Tables (2010–2025), EPA AirData Design Value Reports, and CARB's iADAM Database. *Ozone: No. of Days with High Concentration (>70 ppb) in the year

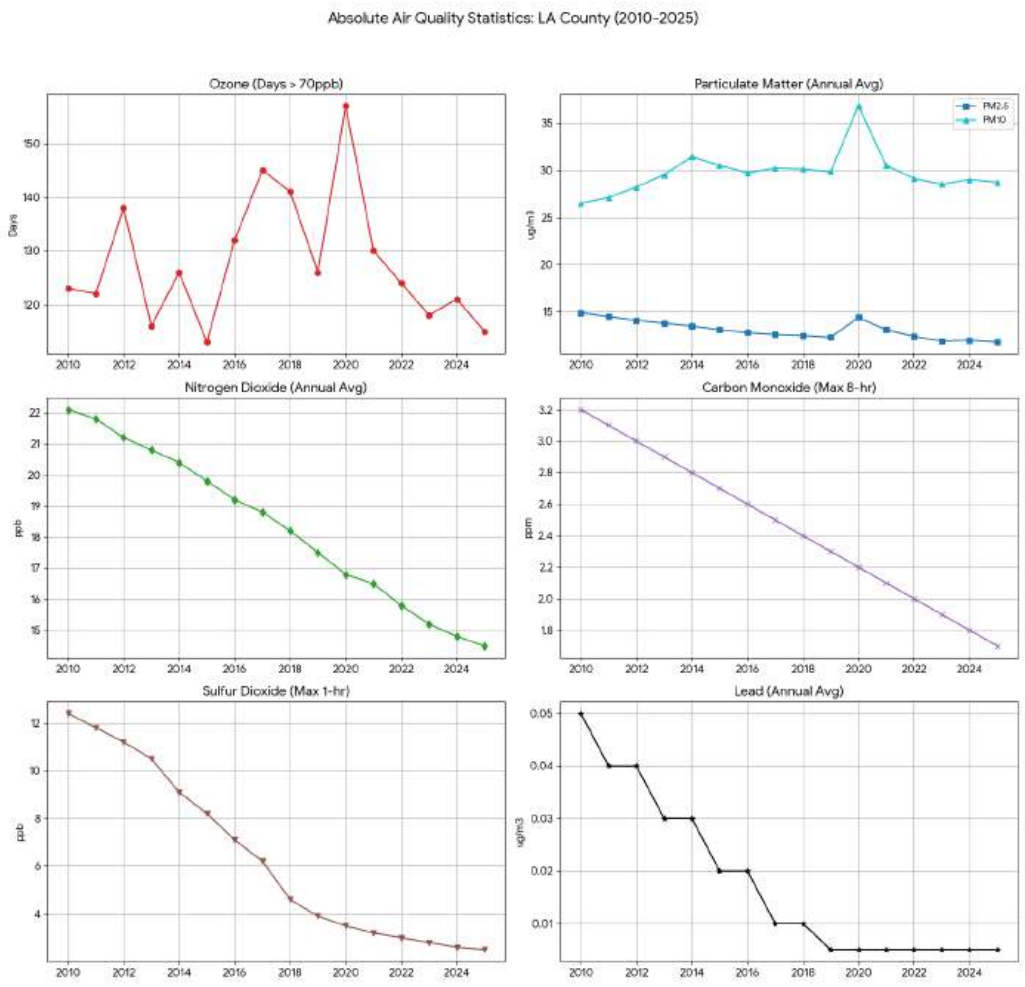


Fig 3: Absolute Air Quality Statistics: LA County (2010-2025). Note: Image generated by Google, 2026 (https://gemini.google.com/).

As shown in Table 2, annual measurements of air pollutant criteria were used to assess regional compliance with state and federal air quality standards in Los Angeles County between 2010 and 2025. Concentrations of O3, NO2, CO, SO2, along with PM2.5, PM10, and Pb were also tracked in the data.

The data represent the spatial averages from air monitoring stations operating within the Los Angeles County boundary. To maintain consistency across the 16-year timeline, annual approximate averages were filtered using regulatory completeness criteria. For cases that are due to environmental disruptions like severe wildfires, they are grouped as exceptional environmental disruptions. They are grouped to exclude emissions of some pollutants caused by temporary irregularities from baseline traffic and industrial sources.

As shown in Table 2 and Figure 3, over the 16-year period, NO2 concentrations dropped by 34.4%, a drop from 22.1 ppb (parts per billion) in 2010 to 14.5 ppb in 2025. At the same time, CO decreased from 3.2 ppm (parts per million) to 1.7 ppm, and the SO2 level decreased by

almost 80%. For Pb, it first dropped significantly, then stabilized below the baseline of <0.01 μg/m^3.

Meanwhile, PM2.5 dropped from 14.9 μg/m^3 to 11.8 μg/m^3. However, in 2020, a rapid and non-linear spike was observed for both PM2.5 and PM10, and PM10 reached a significant 36.8 μg/m^3 in this period. This change is attributable to past wildfire cases rather than tailpipe emissions.

In contrast to the steady decreasing trend demonstrated in both primary vehicle emissions, namely NO2 and CO, secondary photochemical pollutants like O3 still have high volatility. On high ozone days, the level of ozone typically exceeded the 70 ppb safety threshold, and it especially peaks during years with high heat, such as 2017 (145 days) and 2020 (157 days). This indicates that despite success in mitigating primary fossil fuel-related pollutants, variations in climate across regions will still cause secondary atmospheric reactions to proceed.

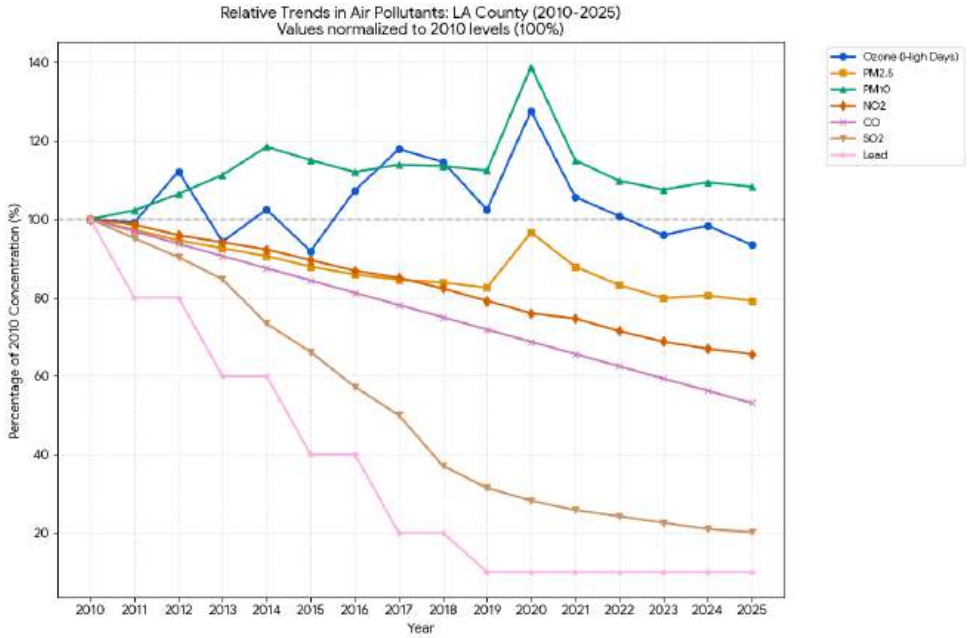


Fig 4: Relative Trends in Air Pollutants: LA County (2010-2025). Note: Image generated by Google, 2026 (https://gemini.google.com/).

With values normalized to the 2010 level (100%), Figure 4 shows the relative trends of each of the six air pollutants in LA County from 2010 to 2025.

Compared to 2010, SO2 and Pb show a significant reduction pattern over the study period, dropping by about 80% and 90%, respectively. This significant reduction suggests the successful elimination of point-source inputs from industrial activities.

NO2 and CO show a steady decreasing pattern. This linear reduction pattern is closely related to previous ZEV deployment rates in LA County.

O3 and PM2.5 concentrations shift away from the standard linear reduction pattern. A significant increase in concentration was observed in both 2017 and 2020. These increases

directly correlate with severe regional wildfires and prolonged heatwaves, highlighting how climate factors and disruptions from local mobile emissions controls temporarily affect concentrations.

For PM10, values exhibit high volatility throughout the 16-year period. This suggests that fuel type change does not affect PM10 much. Past research has shown this pattern is primarily attributed to readjusted road dust, tire wear, and brake dust.

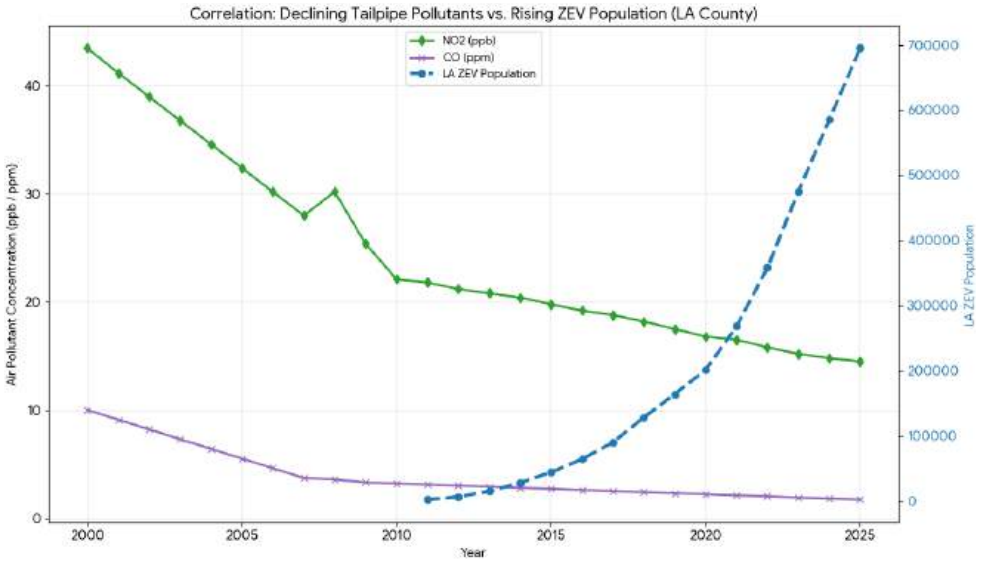


Fig 5: Correlation – Tailpipe Pollutants vs. ZEV Population. Note: Image generated by Google, 2026 (https://gemini.google.com/).

Figure 5 shows the strong correlation between the sudden increase in the ZEV population in Los Angeles County, which started around 2011, and the significant linear increase in NO2 and CO. While NO2 was high above 40 ppb in 2000, it fell to about 14.5 ppb by early 2026.

3. Methodology

3.1 Linear Regression Analysis

To determine whether the increase in the ZEV adoption rate is connected to a decrease in air pollutant concentrations, this study uses an Ordinary Least Squares (OLS) linear regression analysis. Linear regression is a commonly used statistical tool to show the relationship between two continuous variables. (Majka, 2024). In this analysis, the independent variable is time or the level of ZEV adoption, and the dependent variable is air pollutant concentrations.

Mathematical Formulation

1. For each major air pollutant studied in LA County, a simple linear regression model is created to examine the relationship between time (or ZEV adoption rate) and changes in air pollutant concentrations. The model is written as $Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i$, where Y_i

represents a specific pollutant's concentration in a given year, such as the annual average concentration of NO_2, PM2.5, CO, or the number of days with high ozone concentration.

2. X□ represents the independent variable, which can be the year within the 16-year study period (2010–2025) or the number of ZEVs on the road during that year. β_0 is the y-intercept, which represents the estimated starting pollutant level at the beginning of the study period.
3. β_1 is the slope coefficient, which shows the average yearly change in pollutant levels. A negative β_1 value indicates that pollution levels are decreasing over time, while a positive β_1 value shows that the pollution levels are increasing.
4. ε□ represents the error term, which accounts for other factors (weather changes, local emission differences, and natural environmental variations, etc.) that may affect the pollution levels but are not included in the model.

3.2 Estimation and Statistical Testing

The model was calculated using the Ordinary Least Squares (OLS) method, which finds the line that best fits the data by minimizing differences between the actual data and the estimated values. To measure the strength and reliability of the relationship between ZEV adoption rate and air pollutant concentration, this study also uses three statistical measures.

1. The Pearson Correlation Coefficient (r) measures the strength and direction of the relationship between ZEV adoption over time and pollutant levels, with values ranging from -1 (a perfect negative relationship) to +1 (a perfect positive relationship).
2. The coefficient of Determination (R^2) shows the amount of change in pollutant levels that can be explained by the increase in ZEV adoption over time. A higher value suggests the model explains more variation in air quality.
3. The p-value is used to ascertain whether the relationship is statistically significant. A significance level of 0.05 is shown, which means that if the p-value is less than 0.05, the results are considered statistically significant and therefore are unlikely to have occurred by chance.

4. Analysis and Results

To determine whether there are any significant improvements in air quality in LA County from 2010 to 2025, the model was applied independently across seven air pollutants as shown in Table 2: CO, NO2, SO2, Pb, PM2.5, PM10, and ground-level ozone (O3).

Using linear regression enables us to separate tailpipe combustion gases, which demonstrate significant linear reduction due to the retirement of internal combustion engines (ICE). For example, a secondary photochemical compound like O3 and mechanically generated particulates like PM10 exhibit an outlier trend from the original linear reduction paths due to climate variables including extreme heatwaves and road dust readjustment.

The results show a clear division between transportation-related air pollutants and climate-related pollutants. Transportation-related air pollutants have improved significantly,

while climate-related pollutants and local dust, which includes ozone and PM10, have not shown any statistically significant progress.

Table 3. Statistical Significance of Air Quality Trends (2010–2025)

Pollutant	Correlation (r)	P-Value	Statistically Significant?
CO	-1	<0.00001	Yes
NO2	-0.998	<0.00001	Yes
SO2	-0.97	<0.00001	Yes
Pb	-0.92	<0.00001	Yes
PM2.5	-0.804	0.00018	Yes
PM10	0.255	0.34066	No
Ground-level Ozone	-0.032	0.90675	No

A P-value of less than 0.05 indicated that the improvement was statistically significant and unlikely to be due to random chance.

Pollutants most directly correlated with tailpipe emissions (CO, NO, and Pb) showed a significant, near-perfect negative correlation. This dramatic statistical significance implies that shifting to ZEVs and stricter engine standards can mitigate the air chemistry of the LA basin. PM2.5 also showed significant improvement, but it remains unstable due to wildfire smoke.

Even though NOx, a precursor to O3, experienced a significant decrease, the number of unhealthy O3 days still does not show a statistically significant reduction (p=0.906). This trend made O3 especially challenging for LA's air quality. Increasing average temperature works like a catalyst that makes the formation of O3 faster. This characteristic reduced the positive impact made by cleaner cars on the road.

PM10 was the only pollutant that increased slightly. However, this was not statistically significant, which indicates that non-tailpipe emissions such as road dust, tire wear, and construction do not decrease at the same rate as fuel-based emissions. As ZEVs become heavier due to their batteries, this can cause an increase in particles due to tire and road wear. (Jolly 2024)

5. Discussion and Conclusion

This study analyzed air quality trends in Los Angeles County from 2010 to 2025 using Ordinary Least Squares (OLS) linear regression analysis. On top of most existing research that mainly focuses on NO2, our analysis examines a broader range of pollutants, namely CO, NO2, SO2, PM2.5, Pb, O3, and PM10, from LA County over a 16-year period. Our results showed that a statistically significant improvement in pollutants directly correlates with tailpipe emissions.

Over time, CO showed a strong negative correlation with time, which shows a dramatic and continuous decline over the 16-year period. Likewise, NO2, SO2, and Pb showed significant downward trends. PM2.5 also showed a significant improvement, though it still varies, which is likely influenced by periodic wildfires during the study period.

However, not all pollutants followed the same downward pattern. Ground-level ozone (O3), for example, showed no statistically significant change over this period. Instead, it showed a continuous plateau regardless of the reductions observed. It is worth noting that PM10 was the only pollutant that showed a slight increasing trend in the data, even though this increase was not statistically significant.

While the transition toward ZEVs and stricter emissions standards has produced statistically significant improvements in tailpipe-related pollutants (CO, NO2, SO2, PM2.5, Pb), O3 and PM10 levels remain high because they are primarily influenced by climate-related factors and non-tailpipe sources such as road dust and tire wear. The plateau O3 levels imply that climate change may be offsetting the progress made through the shift to ZEVs, while the PM10 trend draws attention to non-exhaust emissions like tire wear and road dust. It is worth noting that EV batteries' heavy weight could also contribute to increased PM10 from both tires and roads, potentially offsetting some of the progress made. This highlights that replacing internal combustion engines with EVs alone is not sufficient to solve all air quality problems. We need a combined approach that addresses both tailpipe and non-tailpipe emissions. How to incorporate climate adaptation strategies into transportation policies to ensure that the progress made through EV adoption is not offset by other factors remains a major challenge for policymakers.

Our findings align with Garcia et al. (2023) and Eckel et al. (2026), who found that ZEV adoption was associated with reductions in NO2 and asthma-related ER visits across California. However, unlike the modeling studies conducted in Taiwan (Li et al. 2016) and by Andre et.al. (2020), our study provides real-world observational evidence from LA County spanning 16 years and examines a broader range of pollutants. Our O3 and PM10 findings fill in the limitations that previous studies did not address. Overall, our findings confirm the air quality benefits of ZEV adoption while revealing that these benefits are not uniform across all pollutants.

It is worth noting that ZEV adoption is not the only reason that tailpipe air pollutants have decreased. Factors such as stricter efficiency requirements for internal combustion engine vehicles also contributed significantly to the reduction.

Even though empirical data confirm statistically significant reductions in transportation-related pollutants, such as NO2 and CO, along with the expansion of ZEV adoption, major structural issues call the effectiveness of long-term policy targets into question. California's mandatory requirement of 100% light-duty ZEV sales by 2035 encounters implementation challenges, as underscored by the recent repeal of the commercial electric trucks mandate following changes in federal Clean Air Act waiver enforcement. These findings highlight that sustainable regional air quality improvements will not only depend on technological changes, but also on piloting a complex federal regulatory framework and the pollutants that are climate-driven.

Despite the billions of dollars in support provided by California's MyFirstEV, HVIP, and CCFR programs, etc., it is certainly not enough to address the huge gap toward the ZEV adoption goal of 100% ZEV sales by 2035. However, it certainly increases the chances of first-time ZEV buyers making up their minds. By making ZEVs more accessible to consumers through these programs, it shows California's commitment to cleaner air, and it is the right direction to go. Admittedly, the efforts will not be as effective if only California does that while most other states choose not to. Similarly, it will not be as effective if only a handful of countries aggressively cut down emissions while most other countries choose not to. As a young generation, it is our responsibility to take over the task and carry it forward.

Acknowledgments

The author acknowledges the use of Grammarly to polish the English grammar, syntax, and overall flow of this manuscript, and Google Gemini to create the figures and to help understand how to use linear regression to obtain and interpret the results.

Works Cited

Andre, Marc, et al. "Diesel, Petrol or Electric Vehicles: What Choices to Improve Urban Air Quality in the Ile-de-France Region? A Simulation Platform and Case Study." Atmospheric Environment, vol. 241, 2020, article 117752. https://doi.org/10.1016/j.atmosenv.2020.117752.

California HVIP, https://californiahvip.org/, Last Accessed August 17, 2026.

Eckel, Sarah P., et al. "Zero-Emissions Vehicle Adoption and Satellite-Measured NO_2 Air Pollution in California, USA, from 2019 to 2023: A Longitudinal Observational Study." The Lancet Planetary Health, vol. 10, no. 1, 2026, article 101379. https://doi.org/10.1016/j.lanplh.2025.101379.

Garcia, Erika, et al. "California's Early Transition to Electric Vehicles: Observed Health and Air Quality Co-Benefits." Science of the Total Environment, vol. 867, 2023, article 161761. https://doi.org/10.1016/j.scitotenv.2023.161761.

Governor Newsom, "Governor Newsom Announces California Will Phase Out Gasoline-Powered Cars & Drastically Reduce Demand for Fossil Fuel in California's Fight Against Climate Change", September 23, 2020, https://www.gov.ca.gov/2020/09/23/governor-newsom-announces-california-will-phase-out-gasoline-powered-cars-drastically-reduce-demand-for-fossil-fuel-in-californias-fight-against-climate-change/

Governor Newsom, "Governor Newsom announces California's new $1 billion rebate program for electric trucks, as Trump cedes global clean vehicle market to China", May 13, 2026, https://www.gov.ca.gov/2026/05/13/governor-newsom-announces-californias-new-1-billion-rebate-program-for-electric-trucks-as-trump-cedes-global-clean-vehicle-market-to-china/

Jolly, Jasper. "Do Electric Cars Have an Air Pollution Problem?" The Guardian, Feb 26, 2024, https://www.theguardian.com/business/2024/feb/26/electric-cars-air-pollution-problem-brakes-tyres

Li, Ning, et al. "Potential Impacts of Electric Vehicles on Air Quality in Taiwan." Science of the Total Environment, vols. 566–567, 2016, pp. 919–928. https://doi.org/10.1016/j.scitotenv.2016.05.105.

Majka, Marcin. (2024). Ordinary Least Squares. September 2024, https://www.researchgate.net/publication/384403324_Ordinary_Least_Squares

Martinez, Anthony. "Governor Newsom Secures 13 Automakers to Offer Instant Rebates for First-Time ZEV Buyers." Governor of California, 16 July 2026, https://www.gov.ca.gov/2026/07/16/governor-newsom-secures-13-automakers-to-offer-instant-rebates-for-first-time-zev-buyers/

Ohnsman, Alan. "California Will Give You $3,500 to Buy a New EV." Forbes, 8 July 2026, https://www.forbes.com/sites/alanohnsman/2026/07/08/california-will-give-you-3500-to-buy-a-new-ev/.

Roy, Jessica. "California Will Pay You to Buy Your First EV; Some Aren't Eligible." San Francisco Chronicle, 26 July 2026, https://www.sfchronicle.com/personal-finance/article/zev-electric-car-rebate-california-22352947.php.

Tefu, Bo. "Starting This Summer, California Car Buyers Can Get an Instant $3500 Off the Cost of Electric Vehicles." LA Focus, 24 July 2026, https://www.lafocusnews.com/california/starting-this-summer-california-car-buyers-can-get-an-instant-3500-off-the-cost-of-electric-vehicles-0e76abdb

Wilkening, Abigail. "Californians Buying Their First EV Can Save $3500 with New Incentive." Car and Driver, 13 July 2026, https://www.caranddriver.com/news/a71921709/california-first-time-ev-buyer-incentive-3500/.

Wu, Jun, et al. "Exposure of PM2.5 and EC from Diesel and Gasoline Vehicles in Communities Near the Ports of Los Angeles and Long Beach, California." Atmospheric Environment, vol. 43, no. 12, 2009, pp. 1962–1971. https://doi.org/10.1016/j.atmosenv.2009.01.009.

SciPy Developers. "scipy.stats.mstats.linregress." SciPy v1.16.0 Manual, https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.mstats.linregress.html. Accessed Aug 1, 2026.

Evaluating the Economic Scalability and Energy Efficiency of Magnetic Confinement Strategies Adopted by Emerging Fusion Startups for Future Commercial Fusion Power Plants By Anahita Dubey

Abstract

This paper evaluates the engineering and economic performance of various fusion confinement approaches with a focus on their scalability, compactness, and fuel cycles to determine the feasibility of future commercial power generation. It begins by mapping the global fusion sector's public and private initiatives, funding patterns, regional contributions, and confinement methods. Through detailed comparative analysis of major confinement methods such as tokamaks, stellarators, field-reversed configurations (FRCs), spherical tokamaks, and magneto-inertial fusion systems, the study highlights key engineering and operational challenges associated with each. Particular emphasis is placed on fuel selection: comparing D-T, D-D, D–He^3, and p-B^{11} fuels, as well as the implications of ignition thresholds, breeding cycles, and neutron radiation management. The findings suggest that compact and spherical tokamak designs are widely regarded as among the more-promising pathways for near-term deployment, while long-term innovation will likely depend on solving challenges related to tritium breeding, stable confinement, and component scalability. The study also outlines future prospects and research directions essential to realize economically viable and scalable fusion power.

1. Introduction

In recent decades, nuclear fusion has gained increasing prominence in the public and private sectors as a potentially long-term, abundant, safe, and sustainable energy resource. In the twelve months leading up to July 2025, fusion companies raised $2.64 billion, a 178% surge from the previous year (Fusion Industry Association). With an increase in funding from a wide range of investors including deep tech venture capital firms such as Breakthrough Energy Ventures, industrial giants such as Siemens Energy, sovereign and quasi-public funds like the European Innovation Council Fund, and a significant rise in public funding by 84% since last year, the fusion private sector is expanding rapidly, with some predicting grid electricity from fusion before 2035 (Fusion Industry Association). Although not inherently dispatchable, fusion would integrate more easily into the grid if it could function as a load-following resource, adjusting output to demand. Scientific interest in nuclear fusion was accelerated by the success of the Manhattan Project, which enabled systematic studies of deuterium–tritium fusion and laid the foundation for thermonuclear research (Chadwick et al.). Following its international declassification in 1958, large-scale programs such as the T-3 Tokamak in the USSR (Smirnov) and the Joint European Torus (JET) (Rebut) began operation. In 2006, a major step was taken in advancing fusion research with the establishment of the International Thermonuclear Experimental Reactor (ITER) in France. The 2010s brought a shift toward private fusion startups such as Helion and Commonwealth Fusion Systems.

Helion's first fusion power plant, Orion, has signed an agreement with Microsoft to provide it with power by 2029. Similarly, CFS's commercial prototype SPARC, is said to generate net electricity by 2027 (Commonwealth Fusion Systems). However, the recent surge in private fusion investment has drawn skepticism from critics who question fusion's economic and environmental viability, as well as its claims about radioactive waste management (Reinders). Moreover, skeptics also cite economic, material, engineering, and grid-integration challenges that raise doubts about fusion's ability to scale competitively against renewables (Nicholas et al.). In light of these challenges, this paper evaluates the economic scalability and energy efficiency of different fusion strategies to identify the most viable pathway toward operational, deployable fusion power.

1.1 Defining Scalability and Efficiency

Scalability, here, is defined by the following parameters:

- Replicability and Modularity: The confinement strategy should be easily replicated and scaled - both into modular versions for other applications like space propulsion, and upward for higher outputs - supported by a mature supply chain for mass-producing key components such as superconducting magnets, first-wall material, and tritium breeding blankets.

- Economies of Scale: With construction of more plants, because of increased standardization, the construction time and costs should be reduced, leading to an overall reduction of cost per unit of electricity produced.

- Regulating and Siting Flexibility: Fusion power plants should be deployed quickly across different regions without encountering drastically different regulatory, safety and environmental barriers.

- Sustainable Fuel Supply: There should be a secure, scalable and sustainable global supply of fuel to support mass-deployment of fusion power plants.

Efficiency, for the purpose of this paper, will be defined using the given parameters:

- Fuel Type: The fuel should have a high cross-section, achieve high reaction rates and present a high energy yield at comparatively low temperatures. Currently, Deuterium-Tritium (D-T) fuel is the most viable option; however, it does have certain drawbacks, to solve which, further research into other fuel sources, especially aneutronic ones, is being conducted. Moreover, the binding curve below shows that fusing light nuclei significantly increases binding energy per nucleon, confirming that fusion - especially with deuterium and tritium - releases more energy than most other reactions, including fission.

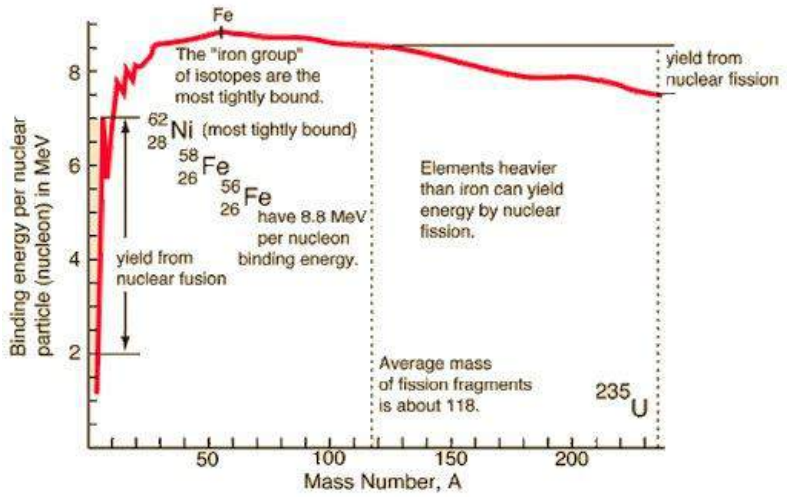


Figure 1. Binding Energy Curve (Nave)

- Confinement Time: Per the Lawson Criterion, net energy production requires a specific combination of plasma density, temperature, and confinement time; a reactor failing to meet these conditions cannot produce an energy surplus. In magnetic confinement fusion, the parameter adjusted to meet the Lawson Criterion is the confinement time. A high confinement time sustains fusion reactions at a lower temperature (preventing instabilities) and a manageable plasma density (reducing collisions and radioactive emissions), enabling net energy gain. The Lawson Criterion is expressed as:

$$n\tau_E \geq \frac{12 k_B T}{\langle \sigma v \rangle Q}$$

Where:

- n = plasma density (particles/m^3)
- τ_E = energy confinement time (s)
- k_B = Boltzmann constant
- T = plasma temperature (K)
- $\langle \sigma v \rangle$ = reactivity (m^3/s)
- Q = energy gain factor

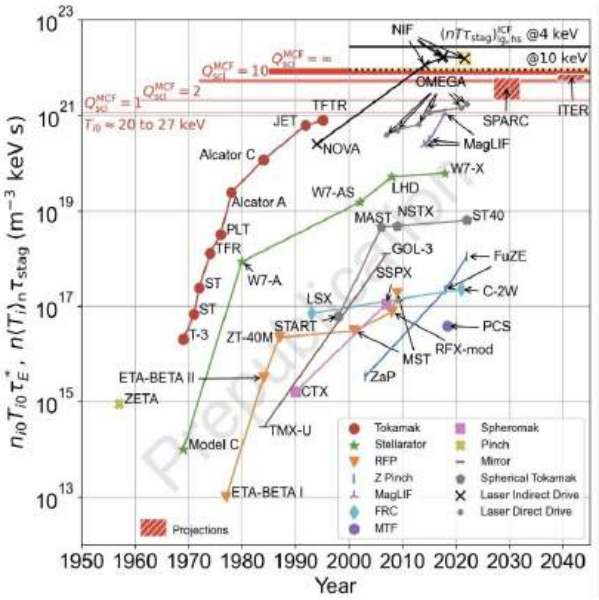


Figure 2. Continuing progress toward fusion energy breakeven and gain as measured against the Lawson criteria (Wurzel and Hsu)

- Size: Compact designs, such as those from Helion and TAE, aim to replace massive infrastructure by achieving high confinement time and net energy gain within a limited plasma volume, increasing energy output per unit volume to enable modular reactors.

- Projected Q Factor: A configuration is efficient if it achieves scientific, engineering, and economic breakeven, each defined in detail in Section 2.2. For a fusion configuration method to be scaled and efficiently employed in future power plants, it must meet the given parameters if fusion is to be pushed to become the primary energy source in the second half of the twenty-first century.

2. Mapping the Fusion Sector: Technologies, Trends, and the Role of Confinement Diversity

2.1. Investment in Diverse Configuration Methods

With more than $2 billion flowing into the sector last year (Fusion Industry Association), fusion research stands at a pivotal point, driven by both scientific promise and the growing energy demands of the AI industry.

Google has entered into the largest direct corporate agreement for fusion energy to date, agreeing to offtake 200 MW of carbon-free electricity from Commonwealth Fusion Systems' first commercial ARC power plant, under development in Chesterfield, Virginia (Gardner, "Google Strikes Deal"). Similarly, Italian energy company Eni, struck a more than $1 billion power purchase with the company's project. The facility is projected to supply approximately 400 MW of clean, zero-carbon power to the electrical grid, thereby advancing the integration of fusion-generated electricity into the energy system (Gardner, "Eni Strikes").

Moreover, Helion Energy has begun construction on nuclear power plants which are planned to be commercially deployable by 2028 to power Microsoft data centers thus not only advancing the timeline to have commercial fusion energy on the grid, but also supporting Microsoft's goal to be carbon negative by 2030 (Nellis). Thus, these developments indicate that tech companies are proactively investing to meet the anticipated energy demands of next-generation data centers which are expected to maintain a compound annual growth rate (CAGR) of 14% through 2029, with the US experiencing a higher growth rate of 18% during the same period (S&P Global Market Intelligence).

Beyond powering data centers, a significant portion of recent private investment is motivated by the diverse confinement methods employed by various companies and their focus on advancing enabling technologies. More than fifty private companies around the world are actively pursuing commercial fusion (Fusion Industry Association) because of the substantial potential payoff. Consequently, investors increasingly favor firms that present clear commercial roadmaps, deploy scalable technologies, and design solutions compatible with existing energy infrastructure. This shift is driving companies to move beyond basic research and toward building deployable, market-ready systems.

Different confinement approaches - from tokamaks and stellarators to magneto-inertial and beam-driven designs - each address specific challenges in achieving sustained, economically viable fusion, broadening the solution space and accelerating innovation. Rather than being redundant, this diversity is mutually reinforcing, allowing cross-pollination toward converged strategies that integrate breakthroughs such as high-temperature superconductors, aneutronic fuels, field-reversed configurations, and advanced materials.

2.2. Projected Q Targets Across Fusion Companies

Q stands for the Fusion Gain Factor, and depending on the context, it represents different thresholds that must be achieved to enable the development of a scalable, commercially viable fusion power plant-regardless of the confinement method employed.

- **Scientific breakeven ($\mathbf{Q}_{scientific} \approx 1$)** refers to the point at which the energy produced by the fusion reactions is equal to the energy delivered to heat the plasma. In 2022, the National Ignition Facility (NIF) at Lawrence Livermore National Laboratory achieved Q = 1.5 using inertial confinement fusion (Abu-Shawareb et al.). However, its scale and energy requirements are unsuited for commercial deployment, and further research is needed to stabilize the reaction. Building on these academic successes, ITER is planning to achieve $Q = 10$ by the mid-2030s and generate 500MW output from 50MW power input (ITER Organization, 2023 Annual Report).

- **Engineering breakeven ($\mathbf{Q}_{engineering} \approx 10$)** occurs when a reactor produces enough energy to operate its internal systems - magnets, cryogenics, and control infrastructure - making $Q \approx 10$ the widely accepted threshold for self-sustaining operation.

- **Electricity breakeven** ($\mathbf{Q}_{electricity} \approx 20 - 30$**)** is the point where a plant generates more electricity than it consumes, accounting for real-world conversion inefficiencies. Most companies prioritize scientific breakeven first, deferring this as an immediate goal (Armstrong et al.). Helion Energy, on the other hand, is pursuing direct electricity generation through field-reversed configuration (FRC), aiming to bypass traditional Q metrics altogether and deliver net electric power (Kirtley and Milroy).

- **Economic breakeven** ($Q_{economic} \approx 100$) marks the stage at which a plant fully recovers costs and turns a profit, shaped by conversion efficiency, construction costs, staffing, uptime, and market prices (International Atomic Energy Agency, Economic Performance Indicators). Since capital costs dominate, companies are developing ways to reduce infrastructure complexity.

2.3. Fusion's Role in Supporting Renewable Grids

To prevent the worst effects of climate change, the global energy grid must transition to low-carbon sources by 2050. However, this transformation should not be framed as an either-or choice between fusion energy and renewables. Rather, the focus should be on developing an integrated grid architecture in which both technologies play complementary roles.

Renewable energy sources such as solar and wind are essential to decarbonization, but they face intrinsic challenges. Their temporal intermittency and low power density can strain power systems, particularly when deployed at scale. Moreover, large-scale renewable deployments require significant land use, which can create constraints especially in highly-populated metropolitans.

Fusion energy offers a compelling solution to many of these challenges. Modular fusion reactors, once operational, could serve as reliable baseload power sources, mitigating the intermittency associated with renewables. Nuclear power is identified as the most land-efficient technology, requiring approximately 50 times less land than coal and 18 to 27 times less than on-ground solar photovoltaic (PV) systems for the same unit of electricity produced (Ritchie), thereby overcoming most constraints of renewable power grids.

While fusion plants require significant upfront capital and operational expenditure, their long-term economics are favorable due to low fuel costs and expenses that decline over time. For an ARC-like plant, CapEx is estimated at approximately €5.42 billion (49% attributable to HTS magnets), with annual OpEx around €113 million - 55% linked to neutron-irradiation component replacements rather than fuel. These replacement costs are expected to fall as materials science improves, and combined with fuel abundance and near-zero variable costs, this points toward a more competitive levelized cost of electricity over the reactor's lifetime (Dusmet Farina and Festa). Fusion also offers safety and environmental advantages over fission: despite larger waste volumes, it contains no transuranic elements, and its waste has lower radioactivity, shorter decay profiles, and greater potential for material recycling (Gonzalez de Vicente et al.).

2.4. Secondary Applications of Fusion

- Heating for Industrial Processes: Fusion's extremely high temperatures can supply industrial heat for energy-intensive industries such as steel, cement, and glass production - a strategically advantageous early market for fusion commercialization (Ward et al.). This heat can also support hydrogen production via high-temperature electrolysis (ITER Organization, "Hydrogen Economy") and enable desalination and chemical refining (Handley et al.) Tokamak Energy has identified industrial heat as one of its primary target markets alongside electricity generation (Fusion Industry Association, Global Fusion Industry 2024).

- Medical Applications: Fusion neutron sources can produce medical isotopes such as Molybdenum-99 (used in over 40 million diagnostic procedures annually), Actinium-225, Copper-67, and Lutetium-177, essential for cancer therapies (National Academies of Sciences, Engineering, and Medicine). Compact fusion devices may also enable Neutron Capture Therapy, a targeted treatment using neutron beams to destroy cancerous cells (Kiyanagi). TAE has developed compact particle accelerators generating focused neutron beams for this purpose, while SHINE Technologies is pursuing beam-target fusion to produce isotopes such as Lutetium-177 (Clark).

- Space Propulsion: Fusion offers a high thrust-to-fuel efficiency ratio and specific impulse, enabling long duration missions to Mars and beyond. In this sphere, aneutronic fuels such as D–He^3 and p–^{11}B are of particular interest, as their charged-particle products enable direct energy recapture for propulsion (Chapman).

- Geothermal Energy Extraction: Traditional geothermal extraction faces geographical limitations, equipment wear, and high costs, with drilling alone accounting for 30-57% of total plant expense (Akindipe and Witter). Quaise Energy is developing millimeter-wave drilling technology - the same waves used to heat fusion plasma - to reach depths of 20 km, unlocking previously inaccessible reservoirs and reducing extraction's environmental footprint (Quaise Energy).

- Scalable Chrysopoeia: Modern chrysopoeia uses nuclear reactions to convert elements like mercury into gold: (n, 2n) reactions, driven by 14 MeV neutrons from D-T fusion plasma, convert mercury-198 into gold-197. This process could yield approximately 2 tons of gold per gigawatt of thermal power annually, potentially doubling a fusion facility's revenue (Rutkowski et al.).

3. Core Challenges in Commercial Fusion Development

3.1. Tritium Breeding Ratio

Among potential fuel cycles for magnetic confinement fusion, the deuterium–tritium (D-T) cycle is the most practical and efficient. With the highest fusion cross-section at the lowest ignition temperatures, D-T enables energy gain under far more accessible plasma conditions than alternatives like DD or D–He^3 (Brereton and Kazimi).

Reactor parameters [a]

Fuel cycle:	DT			DD				DHe
Beta:	5%	10%	20%	5%	10%(HT-9)	10%(RAF)	20%	10%
Major radius (m)	7.6	6.5	5.7	13.0	9.5	9.4	7.3	10.0
Minor radius (m)	1.9	1.6	1.4	3.3	2.4	2.4	1.8	2.5
Inboard blanket/shield [b] (cm)	94	95	96	77	79	80	80	64
Outboard blanket/shield [b] (cm)	205	206	207	181	182	185	184	126
Toroidal field on axis (T)	4.7	3.7	2.9	8.0	7.2	7.1	6.2	7.6
Total wall loading (MW/m^2)	3.8	5.2	6.8	1.1	2.0	1.8	3.3	2.3
Plasma current (MA)	18.4	12.5	8.6	53.8	35.2	34.5	23.2	39.3
Average ion temperature (keV)	11.3	11.3	11.3	33.3	33.3	33.3	33.3	33.3
Fusion power (MW)	3650	3650	3650	2875	2875	2675	2875	3760

[a] All designs produce 1200 MWe with an aspect ratio of 4, an elongation of 2.5, and a safety factor of 2.5.
[b] DT blankets are composed of liquid lithium (71 vol%), helium (22.7 vol%) and HT-9 ferritic steel (6.3 vol%); advanced fuel blankets are composed of helium (20 vol%) and HT-9 (or Fe2Cr1V) (80 vol%). First walls are helium (88 vol%) and HT-9 (or RAF) (12 vol%). Shielding consists of a variable thickness of water cooled Fe1422 (sized for magnet protection and biological shielding), followed by 5 cm of B_4C, and 5 cm of lead (outboard only).

Figure 3. Comparison with different fuel configurations for same reactor parameters (Brereton and Kazimi)

While deuterium is abundant in seawater, tritium is extremely scarce with a global supply of less than 25kg and a half-life of approximately 12.3 years. This is insufficient to fuel even one full-scale fusion power plant for a year. Moreover, there is currently no efficient supply chain for tritium, and the existing tritium from ageing fission reactors, such as the CANDU reactors in Canada, is very limited, with plans for the reactors to be shut down soon (Pearson et al.). Therefore, to achieve successful nuclear fusion and build scalable commercial fusion power plants in the future, it is essential to maintain a Tritium Breeding Ratio (TBR) greater than unity.

$$\text{TBR} = \frac{Number\ of\ tritium\ atoms\ produced}{Number\ of\ tritium\ atoms\ fused}$$

The component associated with TBR in tokamaks is the breeding blanket, which - besides protecting reactor walls from neutron bombardment - reacts neutrons with its material to produce tritium. The breeding blanket is usually made of either solids, liquids (PbLi), molten salts (FLiBe), or pure Lithium. A beryllium (Be) or lead (Pb) neutron multiplier is usually added to increase the number of neutrons and efficiently produce tritium at a higher rate than unity. For the breeder blanket materials, different types of Lithium-based alloys and compounds are proposed as Lithium reacts with neutrons to produce tritium. Reaction of the two stable naturally occurring isotopes of Lithium with neutrons to produce tritium:

$$\text{Li} + \text{n} \rightarrow \alpha + \text{T} + 4.78\ \text{MeV}$$
$$\text{Li} + n \rightarrow \alpha + \text{T} + n + 2.47\ \text{MeV}$$

Molten salts such as FLiBe (LiF and BeF2) remain under research for their thermal

advantages: a high boiling point of 1430°C gives FLiBe the stability to resist melting, decomposing, or reacting at high temperatures. However, the engineering benefits it offers are often offset by its toxicity concerns. Beryllium fluoride is a toxic and carcinogenic compound, small exposures to which can cause incurable lung diseases such as chronic beryllium disease (CBD) and lung cancer. This poses a risk, especially during blanket manufacturing, operation, and maintenance of the reactor.

$PbLi_{100}$ is similarly unfavorable: although its eutectic properties make it easy to store and handle, its high lead content dilutes lithium, making it inefficient for tritium breeding. Moreover, lead is a heavy metal and consequently a slow coolant, due to which it has less heat-transferring capabilities. Not just that, but it can become highly activated under neutron bombardment, leading to an overall high dose rate.

So, in the matter of TBR, the best blanket material appears to be $Li_{17}Pb_{83}$ due to its multiple structural and functional advantages. Unlike FLiBe, which remains under study, $Li_{17}Pb_{83}$ poses fewer toxicity concerns, has a lower melting point (approx. 235°C), and better flow efficiency when properly designed. Moreover, the presence of both breeder and coolant, coupled with greater experimental maturity due to its use in ITER's test blanket modules, makes it a more practical and scalable choice for near-term fusion systems. Its dual functionality in lead enhances useful neutron multiplication, improving tritium breeding when optimally configured.

Liquid breeders also outperform ceramic counterparts, exhibiting higher thermal conductivity and more efficient heat transfer. Additionally, they are easier to fabricate and handle, which simplifies manufacturing. Most importantly, the continuous flow nature of liquid breeders enables inline tritium extraction, whereas tritium extraction from ceramics generally requires thermal cycling or mechanical processing, which can introduce delays and inefficiencies.

In conjunction with these findings, recent studies indicate that lead (Pb) neutron multipliers exhibit superior neutron multiplication performance, particularly at incident neutron energies above 6.5 MeV, due to their comparatively higher neutron interaction cross-section in this energy range, as illustrated in Figure 4 (Goel et al.).

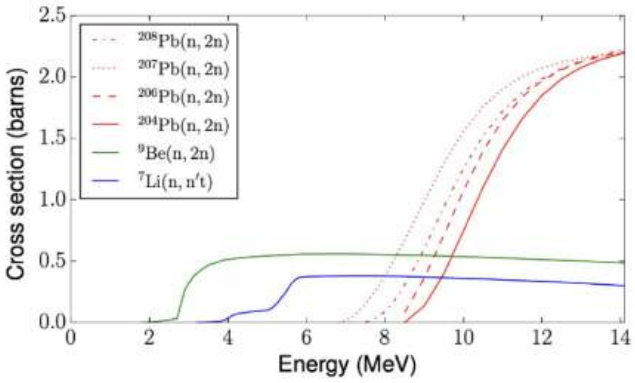


Figure 4. Cross Section vs Energy (Goel et al.)

3.1.1. Investigating the TBR Issue in Diverse Configurations

Stellarators pose a different challenge: while they too use breeding blankets for tritium extraction and cooling, their main issue is relatively poor tritium confinement compared to tokamaks. Unlike tokamaks, which rely on symmetrical toroidal fields generated partly by plasma current, stellarators operate entirely without induced current, using complex external coils to shape their magnetic configuration - making them inherently stable and steady-state capable, but prone to greater neoclassical transport, where particles including tritium drift out of the confined plasma due to magnetic field inhomogeneities. This causes tritium ions to often escape before undergoing fusion, reducing fuel efficiency; the lack of strong shear flow and confinement layering further reduces separation between core and edge regions, compounding premature tritium loss. This inefficiency places greater demand on breeding systems and external fuel supply, significantly affecting the feasibility of a self-sufficient tritium cycle in stellarator-based designs (Xu).

To address the issue of poor tritium confinement in stellarators, recent research has proposed the creation of a three-zone confinement structure within the plasma. In this design, the innermost zone is characterized by poor confinement and is the targeted region for tritium pellet injection, with the expectation that the tritium will subsequently migrate towards the core. However, the physics of pellet deposition remains an area of ongoing investigation and is not yet fully understood. The middle zone serves as the high-confinement core, where deuterium is injected, and fusion will occur. Surrounding this is the outermost edge zone, which is again a low-confinement region and is primarily utilized for radiative heat exhaust and impurity removal. While not yet investigated practically, this concept of structured confinement profile is intended to maximize tritium utilization and maintain plasma purity in stellarator configurations (Boozer).

Renaissance Fusion, a stellarator-based startup, addresses tritium breeding using flowing Li–LiH liquid combined with lead-filled pebbles that act as both neutron multiplier and coolant while shielding the HTS magnets, targeting a TBR above 1.2 without enriched lithium. However, the approach is unproven at scale, and lithium hydride's high reactivity with water and air poses safety challenges (Prost et al.).

Companies such as General Fusion and Tokamak Energy are exploring liquid metal breeding blankets - General Fusion on lead-lithium (Zheng and Todd), and Tokamak Energy via a liquid lithium blanket in its pilot plant (Tokamak Energy).

Other companies bypass tritium breeding entirely: Helion pursues a deuterium–helium-3 cycle where helium-3 is self-generated via D-D fusion (Helion Energy, "Fusion Technology"), while TAE targets aneutronic proton–boron-11 fusion, avoiding tritium fuel-cycle complexities altogether (Magee et al.).

3.2. Stable Confinement

3.2.1. Tokamak

Disruptions are sudden, catastrophic losses of plasma confinement that release vast amounts of energy in short periods and can reduce operational efficiencies as well as damage reactor components. In tokamaks, the poloidal magnetic field is generated by passing an internal current through the plasma. However, this leads to various magnetohydrodynamic (MHD) instabilities, such as kink modes - in which the plasma bends or twists; tearing modes - where the magnetic field lines break and reconnect; and edge-localised modes - formed due to steep gradients as a result of the transition from L-mode to H-mode. These instabilities can trigger vertical displacement events, shifting the plasma column until it touches the reactor walls. This plasma motion induces eddy currents in the reactor walls, causing an abrupt current flow that forms current sheets and dissipates large amounts of heat rapidly. The rapid drop in plasma temperature is known as thermal quench. As the plasma cools, its resistivity increases and current decays; i.e. a phenomenon called current quench takes place - inducing strong electric fields. At a critical field strength, these fields energize electrons enough to strike the reactor wall at high speed, potentially punching through plasma components - carrying MeV energies that make them difficult to dissipate and especially dangerous (Schuller).

Companies like Commonwealth Fusion Systems (CFS), which use tokamaks, have now started building their reactors with high-temperature superconducting (HTS) magnets that produce a stronger magnetic field and thus lead to better confinement as a result of fewer core instabilities (Whyte et al.). CFS is also collaborating with PPPL and MIT's Plasma Science and Fusion Center on machine-learning disruption prediction, using tools such as MIT's DisruptPy library and the Fusion Data Platform (Tang et al.). Furthermore, CFS will incorporate disruption mitigation systems in SPARC such as massive gas injection (MGI) or shattered pellet injection (SPI), which will inject impurities such as neon, argon or deuterium ice to radiate energy uniformly, slow down thermal and current quench and prevent the formation of runaway electrons (Ekmark et al.).

3.2.2. Stellarator

By their very design, stellarators are significantly less susceptible to the kinds of disruptions in tokamaks. This is primarily because stellarators do not rely on strong plasma currents to maintain magnetic confinement, eliminating a major source of magnetohydrodynamic (MHD) instabilities and consequently offering inherently steady-state operation and greater plasma stability (Helander).

Building on these advantages, Renaissance Fusion has engineered its stellarator to eliminate plasma-disruption risk entirely by operating without a plasma current, removing the root cause of current-driven instabilities. Instead, they utilize high-temperature superconducting (HTS) magnets in modular planar coils to generate complex 3D magnetic geometry. If liquid metal enters the plasma - a concern in other configurations - the response is a controlled radiative collapse rather than violent disruption, allowing a rapid, non-destructive restart (Renaissance Fusion).

3.2.3. Field-Reversed Configurations

Lacking a central solenoid, FRCs are compact and modular devices with low magnetic shear as a consequence of the lack of toroidal current in the plasma. This makes them prone to magnetohydrodynamic (MHD) instabilities, such as tilt or shift instabilities, which can alter the orientation of the plasma and cause it to come into contact with the reactor walls. Furthermore, low magnetic shear also increases turbulence-driven transport by failing to disrupt the wave-pattern misalignment that drives instability growth. The cross-field thermal transport of plasma particles also increases due to the lack of a toroidal magnetic field. This leads to an increased radial loss of energy, resulting in a short confinement time (Binderbauer et al.).

To resolve these issues, companies such as TAE Technologies have come up with novel methods. Prolated FRC shapes stabilize the reactor against tilt-mode instabilities by altering magnetic field distribution, while also improving aspect ratio, axial confinement, and pressure balance against external magnetic forces (Romero et al.). TAE also uses magnetic mirrors - stronger fields at the ends than the center - so that, by conservation of energy, particles slow and reverse direction, aiding axial confinement. Another approach undertaken by TAE to improve stable confinement in FRCs is Neutral Beam Injection (NBI). The injected high-energy neutral atoms (hydrogen, deuterium, or tritium) transfer kinetic energy to plasma ions and electrons via collisions, raising temperature and confinement time while damping MHD instabilities and reducing turbulence-driven transport through sheared flow. They also help maintain a more stable configuration by increasing pressure, improving the overall pressure profile (Onofri et al.).

Helion Energy similarly uses a rotating magnetic field to transfer angular momentum to charged particles, inducing azimuthal plasma rotation in its FRC devices. As the plasma rotates, it disrupts the growth of turbulent eddies and introduces sheared flows that stabilize the plasma against micro-instabilities. Through this approach, Helion aims to address challenges inherent in FRC plasmas, such as tilt and rotational instabilities, by leveraging the kinetic nature of the plasma. The resulting high-beta plasma - where the plasma pressure is comparable to the magnetic pressure -enhances confinement and facilitates direct energy recovery from fusion reactions (Helion Energy, "FRC Instabilities").

3.2.4. Spherical Tokamaks

Boasting a compact, high β-configuration, spherical tokamaks can operate at high plasma temperatures; however, this leads to three major microinstabilities: ion temperature gradient modes (ITMs), kinetic ballooning modes (KBMs), and trapped electron modes (TEMs) (Patel et al.). ITMs are drift-wave instabilities formed when the ion temperature gradient becomes too steep, causing heat and ion loss from the core and degrading confinement. KBMs are caused by pressure gradients and magnetic curvature: as β increases in regions of unfavorable curvature, the plasma bulges outward along field lines; because KBMs depend heavily on particle drifts, making them more unstable at lower pressure thresholds, they are studied primarily using gyrokinetic theory models that account for these drifts. TEMs arise from electrons trapped in weak-field regions along their spiral motion - unable to complete a full orbit, these electrons

destabilize the system and raise electron heat losses as they bounce between magnetic mirrors at the top and bottom of the spherical torus, feeling density and temperature pressure gradients; when they resonate with certain wave modes, they transfer energy to the wave, further destabilizing it (Kim and Han).

A potential pathway to mitigate microinstabilities in spherical tokamaks is through optimization of magnetic geometry. For instance, shaping the plasma to achieve negative triangularity can alter the magnetic shear and curvature, thereby reducing the growth rates of ion temperature gradient (ITG) and trapped electron modes (TEM) (Beeke et al.). In addition, because edge instabilities often act as triggers for core microinstabilities, they can be addressed using advanced divertor configurations, which help lower edge temperature and pressure gradients (Havlickova et al.). Tokamak Energy also plans to use high-temperature superconducting (HTS) magnets in its ST25, which offers stronger and more precise magnetic fields; thereby improving magnetic shear and field line control (Sykes et al.). Although theories to control the current density profile in plasma by using radio frequency heating and neutral beam injection to prevent peaking of pressure and temperature are still in the research stage, considerable work is being done on the usage of advanced diagnostics and AI-driven prediction systems to model and predict the microinstabilities (Tang et al.).

3.3. Specific Issues with Scalability from Non-Burning Plasma to Burning Plasma Regime

Today's fusion reactors lie in the non-burning plasma regime, where D-T alpha particles contribute little to heating, so plasma temperature is sustained almost entirely by external sources like RF heating and neutral beam injection. ITER is the first device built to sustain long-duration burning plasma at Q=10 (ITER Organization, "In a Few Lines"), but scaling to this regime poses further issues that must be assessed before selecting a confinement method for future reactors.

3.3.1. Unpredictable Plasma Behaviour

Switching from low-confinement mode to high-confinement mode requires crossing a certain power threshold. However, with an increase in the size and magnetic field strength of the reactor, the power threshold also increases. High-confinement mode offers benefits such as the edge transport barrier, which suppresses turbulence by reducing radial heat and particle transport - though it remains uncertain whether these advantages hold in larger future systems. Most of our current knowledge is from smaller tokamaks, and the findings may not extrapolate linearly due to unpredictable edge-localised modes (ELMs) formed as a result of high-confinement, increased microinstabilities and neoclassical transport effects, and self-heating due to plasma particles, which alters the extremely delicate plasma profile conditions needed to sustain the high-confinement mode (Connor et al.). There is no doubt about the fact that larger devices will experience larger and more damaging ELMs, which are short, intense bursts of particles and heat that erupt from the edge transport barrier. In ITER-class devices, one large ELM can deposit enough energy to melt several millimetres of tungsten in one shot (Paschalidis et al.). An increasingly popular method researched by different companies to solve the issue of large ELMs

is the formation of negative triangularities. High triangularity, shaped using poloidal field coils, lowers the L-mode-to-H-mode power threshold and suppresses edge turbulence, helping avoid large ELMs since the reactor never fully enters H-mode. Negative triangularities are increasingly being focused on for utilisation in conventional and spherical tokamaks (Nelson et al.). For example, the TCV tokamak (Switzerland) showed that negative triangularity plasmas could achieve twice the energy confinement of standard positive triangularity L-mode, which was earlier preferred due to their ability to stabilise MHD instabilities (Coda et al.).

Research led by General Atomics is also being done into Quiescent H-modes, which regulate ELMs by steady oscillations known as Edge Harmonic Oscillation, which carries away excess particles and energy and makes H-mode operation more stable and predictable (Burrell et al.).

3.3.2. Particle and Heat Exhaust

As fusion reactors scale up, the total power output increases much faster than the surface area of the divertor and surrounding walls. This results in extremely high heat fluxes, requiring the divertor to absorb significantly more energy per unit area (Kotschenreuther et al.). In addition, helium nuclei from D-T fusion, once slowed, accumulate as neutral ash; if not removed, this ash displaces fuel and reduces reactivity. It also absorbs thermal energy, further lowering the plasma temperature. Therefore, efficient helium removal becomes critical and is typically achieved by directing plasma flow toward divertors and exhaust systems (Reiter et al.). These systems must remain operational under intense thermal and particle loads, especially as larger reactors target long pulses and steady-state performance.

Because of the high-power exhaust, standard divertor configurations may not withstand this high-power exhaust, making negative triangularity a potential solution. The shape of the triangularity pushes the divertor legs on the outboard side, thus providing a wider surface area for the heat to spread over, lowering flux density at the divertor plates (Kikuchi et al.)

Another approach lies in developing alternative divertor configurations, such as the snowflake divertor, which splits the exhaust over multiple targets, effectively creating multiple strike points on the divertor (Ryutov et al.); liquid metal divertors made using flowing lithium (Nagayama); and radiative divertors where energy is removed from the plasma by radiative cooling, often by introducing impurities (like neon, argon, or nitrogen) that radiate energy in the divertor region (Petrie et al.). Stellarators such as the Wendelstein 7-X employ island divertors, which guide particles to exhaust systems with the help of magnetic islands (Schmitz et al.). Helical divertors are also common in stellarators as they ensure stable, continuous particle flow towards divertor regions (König et al.).

With increased energy exhaust, advanced cooling approaches also need to be developed to sustain long-duration burning plasma operations. The ARC pilot plant by CFS, for example, uses an immersion FLiBe molten salt blanket for efficient heat removal while its high-flux divertor employs actively cooled, high-performance plasma-facing materials to manage extreme heat loads (Kuang et al.). Similarly, Renaissance Fusion employs a liquid lithium–lithium

hydride (Li-LiH) plasma-facing wall in their compact fusion blanket designs, enabling efficient heat extraction for sustained fusion operations (Prost et al.). Furthermore, technologies such as helium gas cooling are being increasingly investigated for fusion due to helium's inert and non-corrosive nature (Baxi and Wong).

3.3.3. Diagnostics and Real-Time Control

As reactor power increases, diagnostic systems must also be improved with better resolution, a fast response rate, and the ability to withstand intense radiation (Pearson and Takeda). As previously discussed, scalability brings about a lot of evolving and unpredictable plasma behaviour, which can be combated with adaptive feedback systems, preferably AI-driven (Churchill). Companies such as TAE have collaborated with Google in the past and collectively worked on improving confinement time by using machine-learning principles to develop the Optometrist Algorithm in TAE's C-2U device, which was the company's advanced beam-driven field-reversed configuration (FRC) experiment at the time (Baltz et al.). Another example of a slightly advanced approach to optimize reactor conditions in specific fuel types include using backpropagation neural networks for autonomous, real-time control of proton-boron ($p\text{-}B^{11}$) fusion reactors, thus addressing the challenges of extreme temperatures and highly nonlinear plasma behavior. The AI system adapts dynamically to changing conditions, aiming to stabilize the plasma and optimize fusion performance, with potential applicability to other fusion systems (Laurelli). Furthermore, MIT's Plasma Science and Fusion Centre has developed an open-source platform called DisruptPy for real-time disruption prediction using neural networks (Trevisan et al.), and the Princeton Plasma Physics Laboratory has collaborated with public-private players such as Commonwealth Fusion Systems (CFS), the U.S. Department of Energy's (DOE), and Oak Ridge National Laboratory (Kremen) and developed a deep neural network called HEAT-ML to find magnetic shadows in a fusion vessel, that is, safe spaces protected from the intense heat of the plasma (Corona et al.).

4. Towards the Ideal Fusion Reactor: A Comparative Analysis of Scalability, Compactness, and Fuel Performance

4.1. Compactness and Scalability

With billions of dollars worth of investment in fusion, capital and operational expenditure remain central points of contention for critics, making compactness, modularity, and scalability key considerations when choosing an ideal confinement method for commercial power plants. Achieving a high-β ratio is comparatively easier in compact reactors because of the larger magnetic field efficiency. Furthermore, they can be efficiently mass produced for integration with industrial facilities, data centres, etc. and installation in populated metropolitan cities as compared to their large-scale renewable counterparts, which occupy a larger land area.

4.1.1. Stellarators

With over two decades taken for the designing and assembly of the Wendelstein 7-X, stellarators are extremely complex fusion configurations with bulky and complex 3-D coils, a low-β configuration and challenges with the TBR (Bosch). Each coil in a stellarator is unique in shape and placement, because of which they must be built accurately to a fraction of a millimetre (Gil et al.). The magnetic coils also experience large electromagnetic forces, requiring huge surrounding cryogenic systems that add further manufacturing hurdles. Stellarators' large size makes scalability uncertain, since their low β-value demands increasingly complex, larger coils to confine large plasma volumes (International Atomic Energy Agency, Stellarator Physics). The three-zone confinement concept remains in research; at present, a single confinement region still causes neoclassical transport losses of tritium, straining breeding systems and fuel supply. Despite this, companies such as Proxima Fusion (Lion et al.) and Type One Energy (Anderson et al.) are planning 2.5 -- 5 m stellarators that avoid current-driven MHD instabilities and achieve steady-state operation with efficient helium exhaust via magnetic islands and helical divertors.

4.1.2. Tokamaks

The main issue with tokamaks lies in their requirement for enormous infrastructure, time-consuming construction, and huge capital expenditure, making siting and scalability a challenge (Gingras and Trépanier). Their reliance on a central solenoid to induce current causes pulsed operation - problematic for baseload power plants, since long-duration pulsed operation remains difficult (Segal et al.). With increasing size, the risk of disruptions such as MHD instabilities, VDEs, and ELMs also increases, becoming harder to predict and control while risking damage to reactor components (Lehnen et al.); mitigating this requires fast, real-time control diagnostics and feedback in modern tokamaks (Peluso et al.). Neoclassical transport effects like banana orbits also grow more significant with size, driving major particle and heat transport (Trinczek et al.), which combined with extreme heat fluxes from power output outpacing exhaust surface area necessitates complex engineering solutions such as liquid metal divertors and advanced cooling systems. Bigger ELMs may pose a greater challenge than predicted, given continued uncertainty around H-mode transition at the higher power thresholds of larger reactors (Peluso et al.). However, companies like CFS are directing more resources toward compact tokamak efficiency, producing ARC (Affordable, Robust, Compact), which uses 23T magnetic fields from HTS magnets to enable a smaller plasma volume while maintaining high fusion power output (Sorbom et al.). This investment stems from tokamaks' rich research history, which has built a strong foundation in plasma confinement physics, disruption handling, and engineering optimization; as size increases, the ratio of power output to energy loss also rises, and plasma self-heating becomes more dominant, increasing the probability of entering the burning plasma regime (Shubov).

4.1.3. Magnetized Targeted Fusion (MTF)

MTF devices are smaller and cheaper than tokamaks or stellarators, using low-cost components like liquid metal liners instead of precision lasers or cryogenic superconductors - but

little is known about the technology at scale. While there are some designs at present, most of them are unproven at scale and need robust testing for commercialization. Furthermore, it is very difficult to achieve scientific breakeven and uniform compression due to not only their pulsed operation but also their smaller plasma volume and lower confinement time. The utilisation of frequent pulsing also leads to extreme mechanical stress and degradation of components (Siemon et al.).

4.1.4. Field Reversed Configurations (FRCs)

With comparatively low energy efficiency and shorter confinement times, FRCs experience high turbulence-driven transport, rapid plasma energy losses and thermal losses due to the lack of magnetic shear and a toroidal magnetic field, which makes it all the more difficult to achieve breakeven. Due to a lack of steady-state or long-duration pulsed operations, they haven't yet demonstrated scalability to the burning plasma regime (Galea et al.). However, FRCs' compact, simple engineering structures make them attractive for space propulsion, enabling light spacecraft with high exhaust velocity and specific impulse - over a hundred times that of chemical propulsion. Since breakeven isn't required for propulsion, small energy losses are tolerable; the goal is converting charged-particle momentum into thrust in short, controlled bursts, with the plasma magnetically shaped to eject only the most energetic ions using simple technology such as magnetic mirrors (Galea et al.).

4.1.5. Spherical Tokamaks

Spherical tokamaks' key advantage is their low aspect ratio, giving them a more spherical than toroidal shape; unlike conventional tokamaks, where the toroidal field dominates, their poloidal field is of comparable intensity. A stronger magnetic field curvature leads to improved confinement for a given field strength, that is, a high-beta value and enhances magnetic shear (Costley and McNamara). Consequently, with an increase in shear, there is suppression of small-scale turbulence and instabilities that are unable to align, which leads to the formation of a steep pressure gradient, that is, the pedestal region and marks the transition to H-mode. Thus, due to the favourable conditions offered by spherical tokamaks to the formation of H-mode, the transition from L-mode to H-mode happens at a comparatively lower power threshold (van Wyk). However, this also increases the probability of ELM formation, to mitigate which devices are investigating methods such as negative triangularity, QH-mode, advanced divertor configurations, and AI-based real-time system controls all of which were discussed in detail in the previous section.

STs also show better scalability as compared to other confinement configurations, as they have an improved performance in smaller volumes and companies like Tokamak Energy are actively working on the rapid iterative development of compact ST reactors with a diameter of about 0.8 m (Windridge).

Compactness has become critical to practical fusion energy, enabling lower costs and easier infrastructure integration. Spherical tokamaks currently offer the most compact, scalable

design, thanks to their low aspect ratio, improved magnetic shear, and lower H-mode power threshold; conventional tokamaks remain the most researched, adapted into high-field versions like SPARC, but still demand large infrastructure to handle instabilities. FRCs, though less efficient for energy generation, are compact and simple enough to suit space propulsion, where specific impulse and size outweigh the need for breakeven.

4.2. Comparative Fuel Performance

The choice of fusion fuel is critical, directly shaping reactor design, material selection, engineering strategy, safety, cost-effectiveness, and scalability (Baker et al.; Brereton and Kazimi). Based on the above considerations, an ideal fuel would possess a low ignition threshold, minimal radioactive byproducts, be abundant and sustainable, readily available, and viable from both economic and technological standpoints.

4.2.1. Deuterium-Tritium (D-T)

Among all fusion fuel options, D–T stands out for its low ignition temperature and highest fusion cross-section, meaning it has the highest probability of undergoing fusion under achievable reactor conditions. This makes D-T the most practical and widely adopted fuel choice in current fusion research and reactor designs (Entler et al.). It is currently being used in ITER's experimental reactor (ITER Organization, "Magic Cocktail"), in Commonwealth Fusion System's SPARC tokamak (Rodríguez-Fernández et al.), and the ST40 by Tokamak Energy (Baylor and Wilson).

However, despite its advantages, D–T fuel presents key challenges, specifically in relation to tritium which is a rare and radioactive isotope. To address this, modern reactor designs plan to implement liquid breeding blankets composed of $Li_{17}Pb_{83}$ with lead (Pb) acting as a neutron multiplier and enhancing tritium production through interactions with lithium, thereby enabling a closed fuel cycle and improving the sustainability of D–T fusion (Moro et al.).

4.2.2. Deuterium-Deuterium (D-D)

D–D fuel is abundant and accessible, removing the need for expensive breeding blanket systems and simplifying reactor design. However, D–D fusion requires almost twice the ignition temperature compared to D–T, and it has a much lower fusion rate, which makes it harder to sustain net-positive energy output (Onofrio).

One unique aspect of D–D fusion is that its reaction produces not just energy, but also tritium and helium-3 as byproducts. These byproducts are valuable - tritium can be recycled back into the fuel cycle, and helium-3 is especially important because of its potential use in advanced aneutronic fusion.

4.2.3. Deuterium-Helium-3 (D-He^3)

D−He^3, an aneutronic fuel, produces minimal neutron radiation, drastically reducing wall damage, shielding needs, and long-term radioactive waste. Additionally, because the reaction

yields mostly charged particles, it allows for the possibility of direct electricity conversion, bypassing the need for large steam turbines. These features make D−He^3 fusion highly appealing for compact, low-maintenance reactors.

However, D−He^3 fusion requires very high ignition temperatures and has a low fusion cross-section (El-Guebaly and Zucchetti). Helion Energy uses a pulsed, non-ignition fusion system with an aim to directly recover clean electricity for which they utilise D−He^3 fusion and source the extremely rare He^3 from the by-products of D-D fusion (Helion Energy, "FAQ").

4.2.4. Proton-Boron-11 (p-B^{11})

p−B^{11} fusion is an aneutronic reaction that produces no neutron radiation under ideal conditions. The reaction proceeds as:

$$p + {}^{11}B \rightarrow 3\ {}^{4}He + 8.7\ MeV$$

This makes it attractive for clean energy, producing only alpha particles that enhance plasma self-heating while minimizing radioactive waste, and enabling direct energy conversion without steam turbines. However, p−B^{11} faces significant challenges: a higher ignition temperature and lower fusion cross-section than D−He^3 result in minimal energy efficiency, while Bremsstrahlung radiation loss cools the plasma faster than it can be heated at high densities, severely limiting net energy gain (McKenzie et al.). TAE Technologies is currently leading p−B^{11} fusion efforts, using an FRC to achieve the necessary confinement and stability (Gota et al.).

Deuterium–Tritium (D–T) remains the most promising fuel for near-term fusion because its low ignition temperature and high fusion cross-section make it far easier to achieve and sustain reactions compared to advanced aneutronic fuels. While alternatives like D–He^3 or p–B^{11} produce less neutron radiation, they demand much higher temperatures and face severe efficiency losses, making them less practical with current technology. Moreover, D–T's main challenge - tritium scarcity - is being actively addressed through breeding blankets, enabling a fuel cycle that supports its long-term viability.

5. Summary and Outlook

This study presented a detailed techno-economic analysis of next-generation fusion energy systems, examining how confinement configuration, fuel type, and reactor design influence viability, cost, and scalability. Among the various methods explored, compact tokamak-based systems - particularly spherical tokamaks - emerged as strong contenders due to their high magnetic efficiency, favorable ratios, and improved confinement in smaller volumes. D–T fusion continues to be the most practical choice in the near term because of its low ignition temperature and high reaction cross-section. Challenges related to tritium scarcity are being addressed through the development of liquid breeding blankets, enabling a self-sustaining fuel cycle and enhancing long-term feasibility.

Crucially, the paper highlights that there is no one-size-fits-all approach to fusion development. Each confinement configuration presents its own trade-offs in compactness,

efficiency, scalability, and fuel compatibility. Therefore, the current wave of global investment-amounting to billions of dollars-across diverse confinement concepts is not only justified but essential. Supporting parallel innovation streams increases the likelihood of breakthrough technologies emerging and reduces the overall risk of dependence on any single approach.

Finally, as the world transitions to cleaner energy systems, fusion must be seen not as a competitor to renewables, but as a foundational complement. Its ability to provide clean, high-density, dispatchable power makes it an ideal partner to intermittent sources like solar and wind. Integration of fusion into renewable grids will be vital for building a resilient, carbon-free energy infrastructure, and future fusion designs must evolve not just for breakeven, but for real-world deployment within dynamic grid ecosystems.

Works Cited

Abu-Shawareb, H., et al. “Achievement of Target Gain Larger than Unity in an Inertial Fusion Experiment.” Physical Review Letters, vol. 132, no. 6, 2024, article 065102.

Akindipe, Dayo, and Erik Witter. “2025 Geothermal Drilling Cost Curves Update.” Proceedings, 50th Workshop on Geothermal Reservoir Engineering, Stanford University, 2025.

Anderson, David T., et al. “A Comprehensive, Unified Baseline Physics Design for the Type One Energy Stellarator Fusion Pilot Power Plant, ‘Infinity Two.’” Journal of Plasma Physics, vol. 91, 2025, article e65.

Armstrong, Robert C., et al. The Role of Fusion Energy in a Decarbonized Electricity System. MIT Energy Initiative, 2024.

Baker, C. C., et al. “Impact of Alternate Fusion Fuels on Fusion Reactor Technology: An Initial Assessment Study.” 8th Symposium on Engineering Problems of Fusion Research, U.S. Department of Energy, 1979.

Baltz, E. A., et al. “Achievement of Sustained Net Plasma Heating in a Fusion Experiment with the Optometrist Algorithm.” *Scientific Reports*, vol. 7, 2017, article 6425.

Baxi, C. B., and C. P. C. Wong. “Review of Helium Cooling for Fusion Reactor Applications.” *Fusion Engineering and Design*, vol. 51-52, 2000, pp. 319–324.

Baylor, L. R., and D. Wilson. “Conceptual Design of a Pellet Injector for the ST40.” *INFUSE*, Oak Ridge National Laboratory, 2022.

Beeke, O., et al. “Impact of Shaping on Microstability in High-Performance Tokamak Plasmas.” *Nuclear Fusion*, vol. 61, no. 6, 2021, article 066020.

Binderbauer, M. W., et al. “A High Performance Field-Reversed Configuration.” *Physics of Plasmas*, vol. 22, no. 5, 2015, article 056110.

Boozer, Allen H. “Stellarators with Enhanced Tritium Confinement and Edge Radiation Control.” *arXiv*, arXiv:2402.08694, 2024.

Bosch, H.-S. “Progress and Challenges in the Construction of Wendelstein 7-X.” *2012 IEEE 24th Symposium on Fusion Engineering (SOFE)*, IEEE, 2012.

Brereton, Sandra J., and Mujid S. Kazimi. “A Comparative Study of the Safety and Economics of Fusion Fuel Cycles.” *Fusion Engineering and Design*, vol. 6, 1988, pp. 207–227.

Burrell, K. H., et al. “Quiescent H-Mode Plasmas in the DIII-D Tokamak.” *Plasma Physics and Controlled Fusion*, vol. 44, no. 5A, 2002, article A253.

Chadwick, M. B., et al. “The Earliest DT Nuclear Fusion Discoveries.” *arXiv*, arXiv:2302.04206, 2023.

Chapman, John J. “Advanced Fusion Reactors for Space Propulsion and Power Systems.” *NASA Technical Reports Server*, NASA Langley Research Center, 2011.

Churchill, R. Michael. “AI Foundation Models for Experimental Fusion Tasks.” *Frontiers in Physics*, vol. 12, 2024, article 1531334.

Clark, Stuart. “Fusion Power Might Be 30 Years Away but We Will Reap Its Benefits Well Before.” *The Guardian*, 11 Aug. 2024.

Coda, S., et al. “Enhanced Confinement in Diverted Negative-Triangularity L-Mode Plasmas in TCV.” *Plasma Physics and Controlled Fusion*, vol. 64, no. 1, 2022, article 014004.

Commonwealth Fusion Systems. “SPARC: Proving Commercial Fusion Energy Is Possible.” *Commonwealth Fusion Systems*. Accessed 18 Aug. 2026.

Connor, J. W., et al. “Edge Localised Modes (ELMs): Experiments and Theory.” *AIP Conference Proceedings*, vol. 1013, no. 1, 2008, pp. 174–190.

Corona, D., et al. “Shadow Masks Predictions in SPARC Tokamak Plasma-Facing Components Using HEAT Code and Machine Learning Methods.” Fusion Engineering and Design, vol. 217, 2025, article 115010.

Costley, A. E., and S. A. M. McNamara. “Fusion Performance of Spherical and Conventional Tokamaks: Implications for Compact Pilot Plants and Reactors.” Plasma Physics and Controlled Fusion, vol. 63, no. 3, 2021, article 035005.

Dusmet Farina, Roberto, and Federica Festa. Assessing the Economic Viability of an ARC-Like Magnetic Fusion Power Plant: Cost Estimation, LCOE, and Sensitivity Analyses. Master’s thesis, Politecnico di Milano, 2023.

Ekmark, I., et al. “Runaway Electron Generation in Disruptions Mitigated by Deuterium and Noble Gas Injection in SPARC.” Journal of Plasma Physics, vol. 91, no. 3, 2025, article E82.

El-Guebaly, L., and M. Zucchetti. “D-3He Fueled Fusion Devices as Step Towards Total Fusion Safety.” 13th International Conference on Emerging Nuclear Energy Systems (ICENES 2007), Istanbul, 2007.

Entler, Slavomir, et al. “Estimation of Fuel Operating Ranges of Fusion Power Plants.” Nuclear Engineering and Technology, vol. 55, no. 7, 2023, pp. 2687–2696.

Fusion Industry Association. The Global Fusion Industry in 2024. Fusion Industry Association, 2024.

Fusion Industry Association. Global Fusion Industry Report 2025. Fusion Industry Association, 2025.

Galea, Christopher, et al. “The Princeton Field-Reversed Configuration for Compact Nuclear Fusion Power Plants.” Journal of Fusion Energy, vol. 42, 2023, article 4.

Gardner, Timothy. “Eni Strikes More Than $1 Billion Power Deal with US Fusion Firm Commonwealth.” Reuters, 22 Sept. 2025.

Gardner, Timothy. “Google Strikes Deal to Buy Fusion Power from MIT Spinoff Commonwealth.” Reuters, 30 June 2025.

Gil, Pedro F., et al. “Manufacturing Tolerances of Non-Planar Coils for an Optimized Tabletop Stellarator.” IEEE Transactions on Applied Superconductivity, vol. 36, no. 3, 2026, p. 1.

Gingras, Yves, and Michel Trépanier. “Constructing a Tokamak: Political, Economic and Technical Factors as Constraints and Resources.” Social Studies of Science, vol. 23, no. 1, 1993, pp. 5–36.

Goel, Vikram, et al. “Optimization of Tritium Breeding Ratio in a DT and DD Submersion Tokamak Fusion Reactor.” arXiv, arXiv:2310.00220, 2023.

Gonzalez de Vicente, Sehila M., et al. “Overview on the Management of Radioactive Waste from Fusion Facilities: ITER, Demonstration Machines and Power Plants.” *Nuclear Fusion*, vol. 62, no. 8, 2022, article 085001.

Gota, H., et al. “Enhanced Plasma Performance in C-2W Advanced Beam-Driven Field-Reversed Configuration Experiments.” *Nuclear Fusion*, vol. 64, no. 11, 2024, article 112014.

Handley, Malcolm C., et al. “Potential Early Markets for Fusion Energy.” *Journal of Fusion Energy*, vol. 40, 2021, article 18.

Havlickova, E., et al. “Investigation of Conventional and Super-X Divertor Configurations of MAST Upgrade Using SOLPS.” *Plasma Physics and Controlled Fusion*, vol. 56, no. 7, 2014, article 075008.

Helander, P. “Stellarator and Tokamak Plasmas: A Comparison.” *Plasma Physics and Controlled Fusion*, vol. 54, no. 12, 2012, article 124009.

Helion Energy. “Frequently Asked Questions.” *Helion Energy*, www.helionenergy.com/faq/.

Helion Energy. “Helion’s Fusion Technology.” *Helion Energy*, www.helionenergy.com/technology.

Helion Energy. “A Note on FRC Instabilities.” *Helion Energy*, 2021, www.helionenergy.com/articles/a-note-on-frc-instabilities/.

International Atomic Energy Agency. *Economic Performance Indicators for Nuclear Power Plants*. IAEA, 2006.

International Atomic Energy Agency. *Stellarator Physics: Proceedings of the Seventh International Workshop, Oak Ridge, 10-14 April 1989*. IAEA-TECDOC-558, IAEA, 1990.

ITER Organization. “The Hydrogen Economy - Brought to You by Nuclear Fusion?” *ITER*, www.iter.org/node/20687/hydrogen-economy-brought-you-nuclear-fusion.

ITER Organization. “In a Few Lines.” *ITER*, www.iter.org/few-lines.

ITER Organization. *ITER 2023 Annual Report.* ITER Organization, 2023.

ITER Organization. “The Magic Cocktail of Deuterium and Tritium.” *ITER*, www.iter.org/node/20687/magic-cocktail-deuterium-and-tritium.

Kikuchi, Mitsuru, et al. “Negative Triangularity as a Possible Tokamak Scenario.” *JPS Conference Proceedings*, vol. 1, 2014, article 015014.

Kim, J. Y., and H. S. Han. “Linear Interaction and Relative Role of the Ion Temperature Gradient and Trapped Electron Modes in the Reactor-Relevant Finite Beta Plasma Condition.” *Physics of Plasmas*, vol. 24, no. 7, 2017, article 072501.

Kirtley, D., and R. Milroy. “Fundamental Scaling of Adiabatic Compression of Field Reversed Configuration Thermonuclear Fusion Plasmas.” *Journal of Fusion Energy*, vol. 42, no. 2, 2023, article 30.

Kiyanagi, Yoshiaki. “Accelerator-Based Neutron Source for Boron Neutron Capture Therapy.” *Therapeutic Radiology and Oncology*, vol. 2, 2018, article 55.

König, R., et al. “Divertors for Helical Devices: Concepts, Plans, Results, and Problems.” *Fusion Science and Technology*, vol. 46, no. 1, 2004, pp. 1–8.
Kotschenreuther, M., et al. “On Heat Loading, Novel Divertors, and Fusion Reactors.” *Physics of Plasmas*, vol. 14, no. 7, 2007, article 072502.
Kremen, Rachel. “Finding the Shadows in a Fusion System Faster with AI.” *Princeton Plasma Physics Laboratory*, 13 Aug. 2025.
Kuang, A. Q., et al. “Conceptual Design Study for Heat Exhaust Management in the ARC Fusion Pilot Plant.” *Fusion Engineering and Design*, vol. 137, 2018, pp. 221–242.
Laurelli, Michele. “AI-Driven Autonomous Control of Proton-Boron Fusion Reactors Using Backpropagation Neural Networks.” *arXiv*, arXiv:2410.12871, 2024.
Lehnen, M., et al. “Disruptions in ITER and Strategies for Their Control and Mitigation.” *Journal of Nuclear Materials*, vol. 463, 2015, pp. 39–48.
Lion, J., et al. “Stellaris: A High-Field Quasi-Isodynamic Stellarator for a Prototypical Fusion Power Plant.” *Fusion Engineering and Design*, vol. 214, 2025, article 114868.
Magee, R. M., et al. “First Measurements of $p^{11}B$ Fusion in a Magnetically Confined Plasma.” *Nature Communications*, vol. 14, 2023, article 955.
McKenzie, Warren, et al. “HB11 - Understanding Hydrogen-Boron Fusion as a New Clean Energy Source.” *Journal of Fusion Energy*, vol. 42, 2023, article 17.
Moro, F., et al. “Nuclear Analyses in Support of the Water-Cooled Lithium Lead Breeding Blanket Design Development: A Prospective Strategy to Achieve the Tritium Self-Sufficiency.” *Fusion Engineering and Design*, vol. 222, 2026, article 115478.
Nagayama, Yoshio. “Liquid Lithium Divertor System for Fusion Reactor.” *Fusion Engineering and Design*, vol. 84, no. 7-11, 2009, pp. 1430–1434.
National Academies of Sciences, Engineering, and Medicine. “Medical Isotope Production and Utilization.” *Molybdenum-99 for Medical Imaging*, National Academies Press, 2016.
Nave, R. “Nuclear Binding Energy.” *HyperPhysics*, Georgia State University, 2010.
Nellis, Stephen. “Helion Energy Starts Construction on Nuclear Fusion Plant to Power Microsoft Data Centers.” *Reuters*, 30 July 2025.
Nelson, A. O., et al. “Robust Avoidance of Edge-Localized Modes Alongside Gradient Formation in the Negative Triangularity Tokamak Edge.” *Physical Review Letters*, vol. 131, no. 19, 2023, article 195101.
Nicholas, T. E. G., et al. “Re-Examining the Role of Nuclear Fusion in a Renewables-Based Energy Mix.” *Energy Policy*, vol. 149, 2021, article 112043.
Onofri, M., et al. “Simulations of Magnetic Field Reversal by Neutral Beam Injection in a Mirror-Confined Plasma.” *Physics of Plasmas*, vol. 32, no. 7, 2025, article 072504.
Onofrio, R. “Concepts for a Deuterium-Deuterium Fusion Reactor.” *Journal of Experimental and Theoretical Physics*, vol. 127, no. 5, 2018, pp. 883–888.
Paschalidis, K., et al. “Impact of Repetitive ELM Transients on ITER Divertor Tungsten Monoblock Top Surfaces.” *Nuclear Fusion*, vol. 64, no. 12, 2024, article 126022.

Patel, B. S., et al. “Linear Gyrokinetic Stability of a High β Non-Inductive Spherical Tokamak.” *Nuclear Fusion*, vol. 62, no. 1, 2022, article 016009.

Pearson, Richard J., and Shutaro Takeda. “Review of Approaches to Fusion Energy.” *Commercialising Fusion Energy*, edited by William J. Nuttall, et al., IOP Publishing, 2020, pp. 2-1–2-31.

Pearson, Richard J., et al. “Tritium Supply and Use: A Key Issue for the Development of Nuclear Fusion Energy.” *Fusion Engineering and Design*, vol. 136, 2018, pp. 1140–1148.

Peluso, Emmanuele, et al. “New Challenges in Nuclear Fusion Reactors: From Data Analysis to Materials and Manufacturing.” *Applied Sciences*, vol. 13, no. 10, 2023, article 6240.

Petrie, T. W., et al. “Application of the Radiating Divertor Approach to Innovative Tokamak Divertor Concepts.” *Journal of Nuclear Materials*, vol. 463, 2015, pp. 1172–1175.

Prost, Victor, et al. “Compact Fusion Blanket Using Plasma Facing Liquid Li-LiH Walls and Pb Pebbles.” *Journal of Nuclear Materials*, vol. 599, 2024, article 155239.

Quaise Energy. “Millimeter Wave Drilling: The Key to Clean Energy Abundance.” *Quaise Energy*, www.quaise.com/news/millimeter-wave-drilling-the-key-to-clean-energy-abundance.

Rebut, Paul-Henri. “The Joint European Torus (JET).” *The European Physical Journal H*, vol. 43, no. 4-5, 2018, pp. 459–497.

Reinders, L. J. *The Fairy Tale of Nuclear Fusion.* Springer, 2021.

Reiter, D., et al. “Helium Removal from Tokamaks.” *Plasma Physics and Controlled Fusion*, vol. 33, no. 13, 1991, pp. 1579–1600.

Renaissance Fusion. “Technology.” *Renaissance Fusion*, renfusion.eu/technology.

Ritchie, Hannah. “How Does the Land Use of Different Electricity Sources Compare?” *Our World in Data*, 2022.

Rodríguez-Fernández, P., et al. “Overview of the SPARC Physics Basis towards the Exploration of Burning-Plasma Regimes in High-Field, Compact Tokamaks.” *Nuclear Fusion*, vol. 62, no. 4, 2022, article 042003.

Romero, J. A., et al. “Magnetic Control of a Field Reversed Configuration.” *Nuclear Fusion*, 2025.

Rutkowski, Adam, et al. “Scalable Chrysopoeia via (n, 2n) Reactions Driven by Deuterium-Tritium Fusion Neutrons.” *arXiv*, arXiv:2507.13461, 2025.

Ryutov, D. D., et al. “A Snowflake Divertor: A Possible Solution to the Power Exhaust Problem for Tokamaks.” *Plasma Physics and Controlled Fusion*, vol. 54, no. 12, 2012, article 124050.

S&P Global Market Intelligence. *451 Research Datacenter Knowledgebase.* S&P Global, 2025.

Schmitz, Oliver, et al. “Stable Heat and Particle Flux Detachment with Efficient Particle Exhaust in the Island Divertor of Wendelstein 7-X.” *Nuclear Fusion*, vol. 61, no. 1, 2021, article 016026.

Schuller, F. C. “Disruptions in Tokamaks.” *Plasma Physics and Controlled Fusion*, vol. 37, 1995, pp. A135–A162.

Segal, D. J., et al. “Steady State versus Pulsed Tokamak Reactors.” *Nuclear Fusion*, vol. 61, no. 4, 2021, article 045001.

Shubov, Mikhail. “Energy Balance Within Thermonuclear Reactors.” *arXiv*, arXiv:2104.06251, 2021.

Siemon, R. E., et al. “Magnetized Target Fusion: Prospects for Low-Cost Fusion Energy.” *Journal of Plasma and Fusion Research SERIES*, vol. 5, 2002, pp. 63–68.

Smirnov, V. P. “Tokamak Foundation in USSR/Russia 1950–1990.” *Nuclear Fusion*, vol. 50, no. 1, 2010, p. 014003.

Sorbom, B. N., et al. “ARC: A Compact, High-Field, Fusion Nuclear Science Facility and Demonstration Power Plant with Demountable Magnets.” *Fusion Engineering and Design*, vol. 100, 2015, pp. 378–405.

Sykes, A., et al. “The ST25 Tokamak for Rapid Technological Development.” 2013 IEEE 25th Symposium on Fusion Engineering (SOFE), IEEE, 2013.

Tang, William, et al. “Implementation of AI/Deep Learning Disruption Predictor into a Plasma Control System.” Contributions to Plasma Physics, vol. 63, no. 3, 2023, article e202200095.

Tokamak Energy. “First Design Details of Fusion Energy Pilot Plant Revealed by Tokamak Energy at Top Physics Conference.” Tokamak Energy, 15 Oct. 2024.

Trevisan, Lucas, et al. “DisruptionPy: An Open-Source Physics-Based Scientific Framework for Disruption Analysis of Fusion Plasmas.” Journal of Open Source Software, vol. 11, no. 119, 2026, article 9364.

Trinczek, Silvia, et al. “Neoclassical Transport in Strong Gradient Regions of Large Aspect Ratio Tokamaks.” Journal of Plasma Physics, vol. 89, no. 3, 2023, article 905890306.

van Wyk, Ferdinand. “Subcritical Turbulence in the Mega Ampere Spherical Tokamak.” arXiv, arXiv:1703.03397, 2017.

Ward, S. H., et al. “Fusion for High-Value Heat Production.” arXiv, arXiv:2504.11072, 2025.

Whyte, D. G., et al. “Experimental Assessment and Model Validation of the SPARC Toroidal Field Model Coil.” IEEE Transactions on Applied Superconductivity, vol. 34, no. 2, 2024, article 0600218.

Windridge, Melanie. “Smaller and Quicker with Spherical Tokamaks and High-Temperature Superconductors.” Philosophical Transactions of the Royal Society A, vol. 377, no. 2141, 2019, article 20170438.

Wurzel, Samuel E., and Scott C. Hsu. “Continuing Progress Toward Fusion Energy Breakeven and Gain as Measured Against the Lawson Criteria.” Physics of Plasmas, vol. 32, no. 11, 2025, article 112106.

Xu, Yuhong. “A General Comparison Between Tokamak and Stellarator Plasmas.” Matter and Radiation at Extremes, vol. 1, no. 4, 2016, pp. 192–200.

Zheng, Shanliang, and Thomas N. Todd. “Study of Impacts on Tritium Breeding Ratio of a Fusion DEMO Reactor.” Fusion Engineering and Design, vol. 98, 2015, pp. 1915–1918.

Evaluating How MATLAB in Python Can Be Used to Plot Planck's 'Black Body Radiation' Spectrum By Akriti Srivastava

Abstract

Planck's formula for 'Black Body Radiation' serves as a fundamental law that used energy quantization to revolutionize modern physics. Plank's law proved that light is not a continuous wave, but is instead composed of packets of energy. This paper would look at how MATLAB in Python can be used to plot Planck's 'Black Body Radiation' spectrum.

Introduction

A blackbody is an idealized object that absorbs any electromagnetic radiation that it comes in contact with. It then emits thermal radiation whose wavelength depends on the temperature of the blackbody.

Planck's law defines the amount of light energy that is given off by a heated object at a certain temperature. The formula for spectral irradiance and the wavelength of the radiation is as follows:

$$E(\lambda, T) = \frac{2hc^2}{\lambda^5} * \frac{1}{e^{\left(\frac{hc}{\lambda kT}\right)} - 1}$$

h = Planck's constant = $6.626 * 10^{-34}\ J * s$
c = speed of light = $2.997925 * 10^{8}\ m / sec$
λ = wavelength (m)
k = Boltzmann's constant = $1.381 * 10^{-23}\ J/K$
T = temperature (K)

Figure 1: *Planck's law [1]*

Method

Plotting Planck's 'Black Body Radiation' spectrum

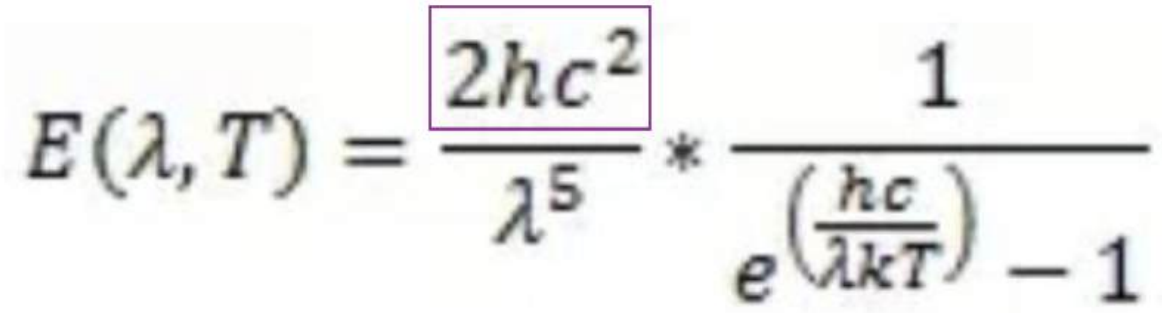
Before any calculations took place, three variables were created for the fundamental physics constants present in the formula for $B(\lambda, T)$. This was done as follows:

```
h = 6.62607015e-34
c = 299792458
kB = 1.380649e-23
```

To perform a standard SI calculation, the wavelength had to be converted from nanometres to meters.

```
#Converting wavelength from nanometers to meters
wavelength_m = wavelength * 1e-9
```

I divided Plank's formula into segments to make the spectral radiance easy to compute. I defined the first part as 'c1'. This includes the first part of the formula, as highlighted below:



$$E(\lambda, T) = \frac{2hc^2}{\lambda^5} * \frac{1}{e^{\left(\frac{hc}{\lambda kT}\right)} - 1}$$

Figure 3: *Planck's formula with 'c1' highlighted*

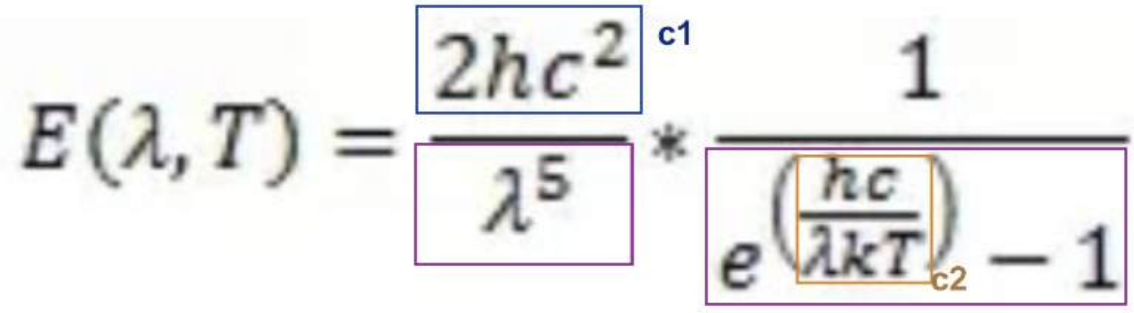
The code for this was as follows:

```
c1 = 2 * h * (c**2)
```

Similarly, I defined 'c2' as the exponent of the natural number 'e'.

```
c2 = (h * c) / (wavelength_m * kB * temp)
```

I then combined c1 and c2 with the remaining parts of the formula, as highlighted in the image below:



$$E(\lambda, T) = \frac{2hc^2}{\lambda^5} * \frac{1}{e^{\left(\frac{hc}{\lambda kT}\right)} - 1}$$

Figure 4: *Planck's formula with remaining region, 'c1' and 'c2' highlighted*

I combined them into a variable for 'Radiance'

```
Radiance = (c1/ (wavelength_m**5)) / (np.exp(c2) - 1)
```

So, the code for the entire function is as follows

```
def Planck(wavelength, temp):

 #Storing required constants
 h = 6.62607015e-34
 c = 299792458
 kB = 1.380649e-23

 #Converting wavelength from nanometers to meters
 wavelength_m = wavelength * 1e-9

 #Performing the calculation
 c1 = 2 * h * (c**2)
 c2 = (h * c) / (wavelength_m * kB * temp)
 Radiance = (c1/ (wavelength_m**5)) / (np.exp(c2) - 1)

 return Radiance * 1e-9
```

I then called this function while plotting my graph. I created a list with the following temperatures: 3000 K, 4000 K, 5000 K, and 6000 K. I also created a list for the following wavelengths: 100 nm, 3000 nm, and 1000 nm. Using a 'for' loop, I called the 'Planck' function and provided it with one wavelength value and one temperature value for these lists as inputs.

```
wavelengths = np.linspace(100, 3000, 1000)
bb_temp = [3000, 4000, 5000, 6000]

for T in bb_temp:
 radiance = Planck(wavelengths, T)
 ax1.plot(wavelengths, radiance, label=f'{T} K', linewidth=2)
```

Once I received my values for radiance, I plotted the curve using matplotlib.pylot.

```
ax1.set_title("Planck Black Body Radiation Spectrum", fontsize = 12,
fontweight='bold')
ax1.set_xlabel("Wavelength (nm)", fontsize = 10)
ax1.set_ylabel("Spectrum Radiation (W/m)", fontsize = 10)
ax1.grid(True, linestyle='--', alpha=0.5)
```

```
ax1.legend(title='Temperature')
```

The graph that I obtained as output was as follows

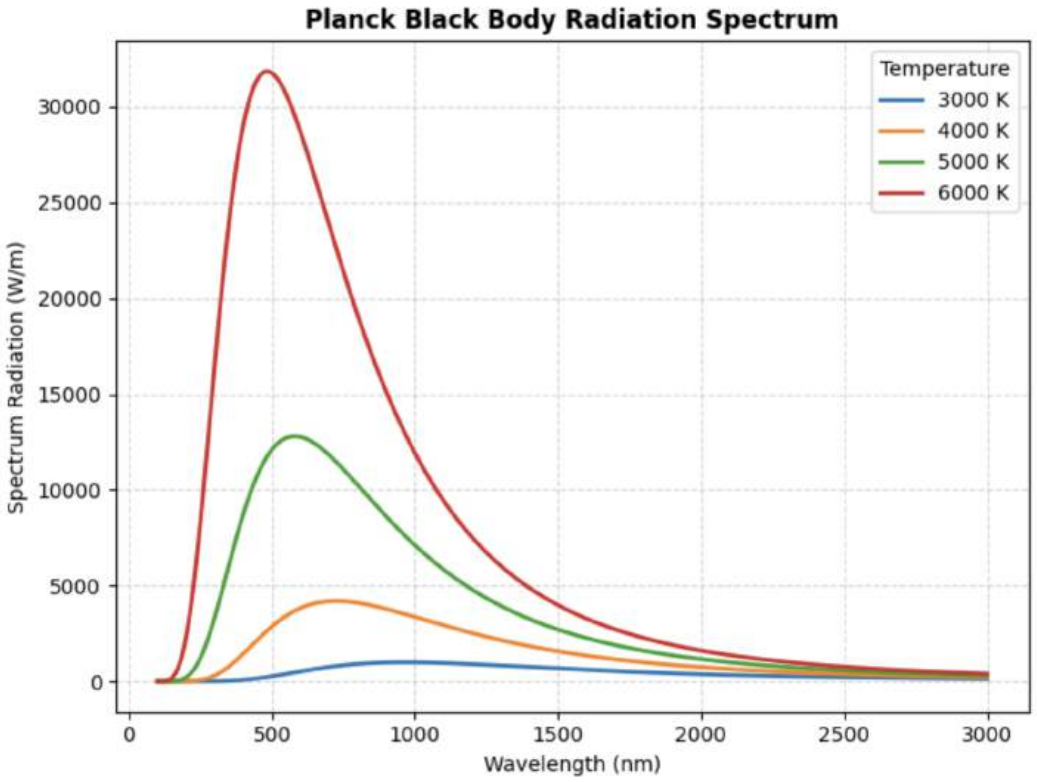


Figure 5: *The graph obtained as an output*

Conclusion

Overall, this paper shows that MATLAB in Python can successfully be used to plot Planck's Black Body Spectrum. The graph obtained as output for Planck's spectrum is correct as it relates to the expected trend of the formula. At small wavelengths, the exponent ('c2') becomes incredibly large, so the exponential term increases. As a result, the spectral radiance is smaller at lower wavelengths. Also, the spectral radiance at peak wavelength at higher temperatures is higher than that for lower temperatures.

Works Cited

[1] Libretexts. “1.1.3: Quantization of Energy And Bohr Model of The Atom.” Chemistry LibreTexts, 18 Jan. 2025, chem.libretexts.org/Courses/Ursinus_College/CHEM322%3A_Inorganic_Chemistry/01%3A_Atomic_Structure/1.01%3A_Historical_Development_of_Atomic_Theory/1.1.03%3A_Quantization_of_Energy_and_Bohr_Model_of_the_Atom.

[2] Guo, Xiaoyang, and Xinzhi Xie. “Planck Formula for Black-body Radiation: Derivation and Applications.” Highlights in Science, Engineering and Technology: PSETE 2023, vol. 64, pdfs.semanticscholar.org/8356/63689a2139503c3914439635effd9bb68596.pdf.

[3] Sanbha Khongngain, Lemuel. “Revisiting the Significance of Planck’s Constant in Quantum Mechanics.” International Journal of Scientific Research in Physics and Applied Sciences, Dec. 2023, www.isroset.org/pub_paper/IJSRPAS/2-ISROSET-IJSRPAS-09306.pdf.

A Comparison Between Direct Heating and Temperature Control Using a Water Bath
By Dahyuck Im

Abstract

In this paper, we compare the temperature changes in direct heating and water bath heating using simple heat-transfer models. Direct heating is modeled as a one-step heat transfer process, where the heat source heats the water directly. Water bath heating is modeled as a two-step heat transfer process, where the heat source first heats the outer water and then the heat is transferred to the inner water. To describe these situations, we define the temperatures, masses, and heat exchange constants needed for the models, and derive the differential equations using the conservation of energy. In the direct heating case, we obtain one differential equation. In the water bath heating case, we obtain a system of differential equations in which two temperature variables are connected to each other. In particular, we write the water bath system in matrix form and analyze the structure of its solution using diagonalization.

1.Introduction

In scientific experiments and cooking processes, it is often important to keep the temperature stable and controlled. For example, chocolate may need to be melted slowly, egg-based sauces may need to be heated carefully, and some laboratory samples may need to be warmed gradually. In these situations, if the temperature rises too quickly, the material may burn, harden, or change its physical or chemical properties. Therefore, it is important to control temperature changes smoothly, rather than simply heating an object as quickly as possible. Figure 1 shows an example of melting chocolate by water bath heating. Since chocolate can burn when it is heated directly, water bath heating is often used.



Fig1:An example of melting chocolate by placing a bowl over heated water. Image from The Mom 100.

Direct heating is a method in which a heat source transfers heat directly to a container or an object. In this case, heat can be transferred quickly, so the temperature may change greatly in

a short time. Also, if the heat is not distributed evenly, one part of the object may reach a much higher temperature than the other parts. This can cause unwanted changes in the material. On the other hand, water bath heating is an indirect heating method. In this method, water is placed in a large pot, and a smaller container is placed inside it. Heat is first transferred to the outer water, and then the water transfers heat to the inner container and its contents. Therefore, since heat is transferred through two steps in water bath heating, we can intuitively expect the temperature change to be more gradual.

In this paper, we mathematically compare the temperature changes in direct heating and water bath heating. Direct heating is modeled as a one-step heat transfer process, where the heat source transfers heat directly to the object. On the other hand, water bath heating is modeled as a two-step heat transfer process, where the heat source first heats the outer water, and then the heat is transferred to the inner container and its contents. After that, we derive and solve the differential equations for each case, and analyze how the temperature changes differently over time in the two methods. Finally, the goal of this paper is to explain mathematically why water bath heating is more useful than direct heating for controlling temperature. These two situations can be represented as follows.

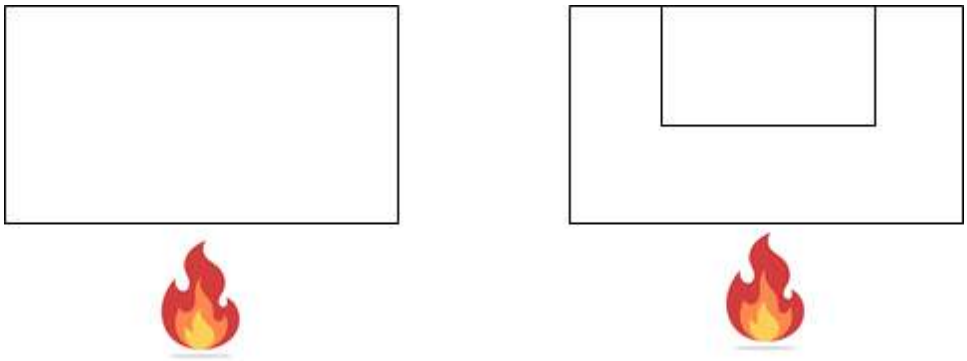


Fig2: Direct heating and water bath heating. The left figure shows direct heating, where the container is heated directly by the flame. The right figure shows water bath heating, where heat is first transferred to the outer water and then to the inner container.

The left figure shows direct heating, and the right figure shows water bath heating. The variables and physical settings for the above models will be introduced in Section 2. This paper is organized as follows. First, in Section 2, we express the given situations using variables and explain what each variable means. In Section 3, we derive the differential equations satisfied by these variables. In Section 4, we introduce a simple method for solving the given differential equations. Finally, in Section 5, we discuss how the results obtained from the differential equations can be applied to real situations.

2.Variables and physical setting

In Section 2, we introduce the variables needed to describe the two heating situations shown in Figure 2. In the case of direct heating, heat is transferred directly from the heat source

to the object, so we build the model mainly using the temperature of the object. On the other hand, in water bath heating, heat is first transferred to the outer water, and then it is transferred to the inner container and its contents. Therefore, we need to use separate variables for the temperature of the outer water and the temperature of the inner contents. In this section, we define these variables and organize the physical setting so that the differential equations can be derived in Section 3.

The mass of the pots is assumed to be very small compared to the total mass of the water. Also, the amount of heat absorbed by the pots is assumed to be small compared to the amount of heat absorbed by the water. Therefore, in this paper, we ignore the heat absorbed by the pots themselves and focus only on the temperature changes of the water.

We also assume that convection occurs well inside each pot, so the temperature of the water is uniform in each region. In the left situation in Figure 2, the whole water inside the pot is assumed to have the same temperature. In the right situation in Figure 2, the water inside the small pot is assumed to have a uniform temperature, and the water between the large pot and the small pot is also assumed to have a uniform temperature. Therefore, we define the variables as follows.

T : the temperature of the water in the direct heating case
T_A : the temperature of the outer water in the water bath heating case
T_B : the temperature of the inner water in the water bath heating case

We denote the specific heat capacity of water by c. The actual value of the specific heat capacity of water is approximately

$c = 4184$ J/kg ・℃ .

However, in this paper, c is treated as a constant. Also, in many later calculations, it will be canceled out, so it does not make the calculation much more complicated.

We also define the mass of water in each region. In the right situation, let M_A be the mass of the water between the large pot and the small pot, and let M_B be the mass of the water inside the small pot. Then, in the left situation, the mass of the water inside the pot can be written as $M_A + M_B$. This means that the left situation can be considered as the combined version of the two water regions in the right situation. Since we ignore the mass of the pots, the total mass is expressed only by the mass of the water.

We also make the following assumption about the heat source. Suppose that the heat source is attached to the bottom of the pot, and the only thing we can control is whether the heat source is turned on or off. When the heat source is turned on, we assume that energy is transferred to the pot at a rate of W per unit time. On the other hand, when the heat source is turned off, no energy is transferred from the heat source to the pot, so the rate of energy transfer is 0. Therefore, in this model, the heat source supplies energy at a constant rate when it is turned

on, and supplies no energy when it is turned off.

After the heat source is turned off, the pot cools down by exchanging heat with the surrounding air. To describe this, we assume that the temperature of the surrounding air is a constant T_0. For example, T_0 can be thought of as room temperature, which is about $23℃$.

Finally, we denote the surface areas related to heat exchange by k_1 and k_2. Here, k_1 represents the surface area of the large pot, and k_2 represents the surface area of the small pot. Then the heat exchange between the outside air and the large pot is proportional to k_1, and the heat exchange between the large pot and the small pot is proportional to k_2. Since the large pot usually has a larger surface area than the small pot, it is natural to assume that $k_1 > k_2$.

If we also consider the lid part of the pots in Figure 2, the calculation becomes much more complicated because there are more paths for heat exchange. Therefore, in this paper, we ignore the heat exchange through the lid. Since the main heat exchange paths are already included, this assumption simplifies the calculation without changing the main physical idea too much.

If we mark the variables defined above on the diagram in Figure 2, we obtain the following figure.

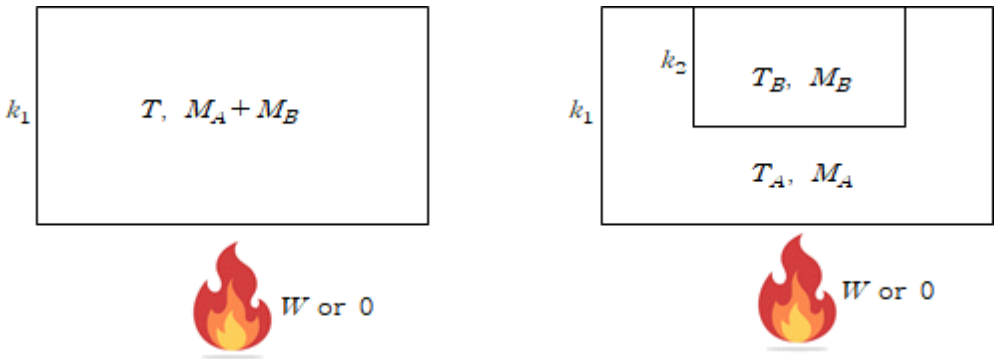


Fig3: Temperatures, masses, and heat exchange areas used in the mathematical models for direct heating and water bath heating.

3.Derivation of the differential equations

In the previous section, we defined the variables needed to describe direct heating and water bath heating mathematically. Now, using these variables, we will construct the differential equations that describe how the temperature changes over time in the two situations. In the case of direct heating, we consider the situation where the heat source transfers energy directly to the water. In the case of water bath heating, the heat source first transfers energy to the outer water, and then this energy is transferred to the inner water. Therefore, the heat transfer structures of the two cases are different, and this difference leads to different differential equations. First, Section 3.1 constructs the differential equation for direct heating, and Section 3.2 constructs the differential equations for water bath heating.

3.1 Differential equation for direct heating

First, we consider the case of direct heating. In direct heating, the heat source transfers heat directly to the water. Let T be the temperature of the water, $M_A + M_B$ be the mass of the water, and c be the specific heat capacity of water.

When the heat source is turned on, energy is transferred to the water at a rate of W per unit time. At the same time, the water exchanges heat with the surrounding air. We assume that the heat lost to the outside is proportional to the temperature difference $T - T_0$ and to the surface area constant k_1. Therefore, the total energy gained by the water per unit time is

$W - k_1(T - T_0)$.

Therefore, when the heat source is turned on, the differential equation for direct heating is

$(M_A + M_B)c\frac{dT}{dt} = W - k_1(T - T_0)$ or equivalently $\frac{dT}{dt} = \frac{W - k_1(T - T_0)}{(M_A + M_B)c}$.

When the heat source is turned off, no energy enters the water from the heat source. Therefore, we substitute $W = 0$ into the equation. Then

$\frac{dT}{dt} = -\frac{k_1(T - T_0)}{(M_A + M_B)c}$.

To write this equation more simply, let

$C = \frac{k_1}{(M_A + M_B)c}$ and $\widetilde{T_0} = T_0 + \frac{W}{k_1}$.

Then, both differential equations (turned on/off cases) can be simplified as follows:

$\frac{dT}{dt} = -C(T - \widetilde{T_0})$ and $\frac{dT}{dt} = -C(T - T_0)$.

Therefore, in the direct heating case, $T(t)$ converges exponentially to T_0 when the heat source is turned off, and $T(t)$ converges exponentially to $\widetilde{T_0}$ when the heat source is turned on.

3.2 Differential equations for water bath heating

Now, we consider the case of water bath heating. In water bath heating, the heat source does not directly heat the water inside the small pot. Instead, the heat source first heats the outer water between the large pot and the small pot, and then the outer water transfers heat to the water inside the small pot. Therefore, the temperature change in water bath heating cannot be described

by only one temperature variable. We need to consider both the temperature T_A of the outer water and the temperature T_B of the inner water.

We represent the heat exchange with the outside air by $k_1(T_A - T_0)$, and the heat exchange between the outer water and the inner water by $k_2(T_A - T_B)$. Therefore, when the heat source is turned on, the temperature changes in water bath heating are described by the following system:

$$M_A c \frac{dT_A}{dt} = W - k_1(T_A - T_0) - k_2(T_A - T_B)$$
$$M_B c \frac{dT_B}{dt} = k_2(T_A - T_B)$$

The first equation describes the temperature change of the outer water. The outer water receives energy from the heat source, but loses energy to the outside air and to the inner water. The second equation describes the temperature change of the inner water. The inner water gains or loses heat depending on the temperature difference between the outer water and the inner water.

First, divide each equation by $M_A c$ and $M_B c$. After that, we express the system differential equations in terms of $T_A - T_0$ and $T_B - T_0$ as follows:

$$\frac{d}{dt}(T_A - T_0) = \frac{W}{M_A c} - \left(\frac{k_1}{M_A c} + \frac{k_2}{M_A c}\right)(T_A - T_0) + \frac{k_2}{M_A c}(T_B - T_0)$$
$$\frac{d}{dt}(T_B - T_0) = \frac{k_2}{M_B c}(T_A - T_0) - \frac{k_2}{M_B c}(T_B - T_0)$$

When the heat source is turned off, we again substitute $W = 0$ to get

$$\frac{d}{dt}(T_A - T_0) = -\left(\frac{k_1}{M_A c} + \frac{k_2}{M_A c}\right)(T_A - T_0) + \frac{k_2}{M_A c}(T_B - T_0)$$
$$\frac{d}{dt}(T_B - T_0) = \frac{k_2}{M_B c}(T_A - T_0) - \frac{k_2}{M_B c}(T_B - T_0)$$

We will investigate this system in the next section.

4. Solution of the differential equations

In Section 4, we solve the differential equations constructed in the previous section. In the case of direct heating, the differential equation can be solved easily using a basic exponential function, and its solution is already well known. Therefore, instead of discussing the solution for direct heating in detail, this paper focuses on the case of water bath heating, where the two

temperature variables T_A and T_B appear together. In particular, we solve the system of differential equations for water bath heating and analyze how the temperatures of the outer water and the inner water change over time.

To write the system for water bath heating more simply, let

$T_1 = T_A - T_0$ and $T_2 = T_B - T_0$.

Then, we have

$$\frac{dT_1}{dt} = \frac{W}{M_A c} - \left(\frac{k_1}{M_A c} + \frac{k_2}{M_A c}\right) T_1 + \frac{k_2}{M_A c} T_2$$
$$\frac{d}{dt} T_2 = \frac{k_2}{M_B c} T_1 - \frac{k_2}{M_B c} T_2$$

To simplify the system, we need to be careful about the sign. To remove the term involving W, it is natural to define

$\widetilde{T_1} = T_1 - \frac{W}{k_1}$ and $\widetilde{T_2} = T_2 - \frac{W}{k_1}$.

If we substitute them into the system of differential equations of T_1 and T_2, then we get

$$\frac{d\widetilde{T_1}}{dt} = -\left(\frac{k_1}{M_A c} + \frac{k_2}{M_A c}\right) \widetilde{T_1} + \frac{k_2}{M_A c} \widetilde{T_2}$$
$$\frac{d}{dt} \widetilde{T_2} = \frac{k_2}{M_B c} \widetilde{T_1} - \frac{k_2}{M_B c} \widetilde{T_2}$$

Now, we write the above system in matrix form. If we collect the two variables into one vector, we can write

$\vec{\widetilde{T}} = \begin{bmatrix} \widetilde{T_A} \\ \widetilde{T_B} \end{bmatrix}$ and $\frac{d}{dt} \vec{\widetilde{T}} = M \vec{\widetilde{T}}$,

where $M = \begin{bmatrix} -\frac{k_1 + k_2}{M_A c} & \frac{k_2}{M_A c} \\ \frac{k_2}{M_B c} & -\frac{k_2}{M_B c} \end{bmatrix}$ is a 2×2 matrix. If we solve this matrix valued differential equation, we get

$$\vec{\widetilde{T}}(t) = e^{Mt} \vec{\widetilde{T}}(0) = P e^{Dt} P^{-1} \vec{\widetilde{T}}(0)$$

when the diagonalization of M is $M = PDP^{-1}$. From some calculations, we can find its diagonalization as follows (Lay et al.):

$$P = \begin{bmatrix} \frac{k_2}{M_A c} & \frac{k_2}{M_A c} \\ \frac{k_1+k_2}{M_A c}+\lambda_1 & \frac{k_1+k_2}{M_A c}+\lambda_2 \end{bmatrix} \text{ and } D = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \text{ where}$$

$$\lambda_1 = -\frac{\frac{k_1+k_2}{M_A c}+\frac{k_2}{M_B c}-\sqrt{\left(\frac{k_1+k_2}{M_A c}-\frac{k_2}{M_B c}\right)^2+4\left(\frac{k_2}{M_A c}\right)\left(\frac{k_2}{M_B c}\right)}}{2} \text{ and}$$

$$\lambda_2 = -\frac{\frac{k_1+k_2}{M_A c}+\frac{k_2}{M_B c}+\sqrt{\left(\frac{k_1+k_2}{M_A c}-\frac{k_2}{M_B c}\right)^2+4\left(\frac{k_2}{M_A c}\right)\left(\frac{k_2}{M_B c}\right)}}{2}.$$

Finally, if we get the following solution:

$$\widetilde{T_1}(t) = C_1 e^{\lambda_1 t} + C_2 e^{\lambda_2 t} \text{ and } \widetilde{T_2}(t) = C_1\left(\frac{\frac{k_1+k_2}{M_A c}+\lambda_1}{\beta}\right)e^{\lambda_1 t} + C_2\left(\frac{\frac{k_1+k_2}{M_A c}+\lambda_2}{\beta}\right)e^{\lambda_2 t}.$$

Since C_1 and C_2 can be determined by initial temperatures $\widetilde{T_1}(0)$ and $\widetilde{T_2}(0)$, we can fully solve $\widetilde{T_1}(t)$ and $\widetilde{T_2}(t)$, although the expression is really complicated.

5. Application and interpretation of the results

Finally, using the results obtained above, we can predict the temperature change of the inner contents during water bath heating. When the heat source is turned on, the initial temperatures and the physical constants can be used to calculate the temperature of the inner contents after a certain amount of time. In other words, if we know the initial temperatures of the water in the large pot and the small pot, we can estimate the temperature of the inner contents as time passes.

A similar method can also be used when the heat source is turned off. In this case, we assume that no energy enters from the heat source, and the initial temperatures can be used to calculate how the temperatures of the outer water and the inner water change over time. Therefore, this model allows us to estimate when the inner contents will reach the desired temperature, even without using a thermometer at every moment. It also helps us decide when to turn the heat source on and off so that the contents do not become too hot. In this sense, the model shows that water bath heating is not only an empirical cooking method, but also a process that can be understood mathematically as a way to control temperature changes.

Works Cited

Lay, David C., Steven R. Lay, and Judi J. McDonald. Linear Algebra and Its Applications. 5th ed., Pearson, 2016.

The Mom 100. “How to Make a Double Boiler.” The Mom 100, https://themom100.com/how-to-make-a-double-boiler/.

Beyond Chance and Necessity: Cosmic Fine-Tuning at the Boundary of Cosmology and Teleology By Camila Richardson

Abstract

The physical universe depends on a striking range of conditions. Stars require particular conditions to form and persist, nuclear reactions depend upon relationships among fundamental forces, and the large-scale structure of the universe is influenced by the value of the cosmological constant. Alter some of these conditions, and galaxies, stars, stable chemistry, and the structures associated with life may never emerge. This apparent sensitivity is commonly described as cosmic fine-tuning.

The physical observation, however, does not settle the question of what fine-tuning means. A life-permitting universe can be interpreted in several ways. Multiverse proposals suggest that observers necessarily find themselves in one of many physically possible environments. Teleological arguments interpret the same order as evidence of intentional selection. Other approaches appeal to a deeper physical theory that would show why the observed parameters could not have been otherwise, while a brute-fact position treats their values as fundamental. These explanations rest on different assumptions and answer different questions.

The central difficulty is the movement from observation to explanation. Demonstrating that physical conditions are sensitive does not by itself establish that those conditions are improbable, intentionally selected, or produced within a multiverse. Probability requires a defined set of alternatives and a method for assigning probabilities to them. Multiverse proposals require a physically motivated ensemble and a defensible measure, while teleological arguments require philosophical premises that extend beyond physics. Fine-tuning therefore occupies an unusual position between cosmology and philosophy. Physics can determine how strongly the emergence of complex structure depends on the universe's parameters; the deeper question is what, if anything, that dependence tells us about why the universe has the structure it does.

Introduction: The Question Behind the Constants

The equations of physics contain quantities whose values are measured with remarkable precision but whose ultimate origins remain unclear. The masses of elementary particles, the strengths of fundamental forces, and the cosmological constant all enter our descriptions of nature, yet current theories do not provide a complete account of why these particular values characterize our universe. Their significance becomes especially apparent when we consider counterfactual possibilities. Alter certain parameters, and atoms may become unstable, stars may fail to form or persist, or the chemistry required for complex structures may change dramatically.

The universe is therefore sensitive to the values of its fundamental parameters. Those values influence which structures can exist and how the universe develops over time. They affect the stability of atoms, the formation and longevity of stars, the production of heavy elements, and the development of increasingly complex structures. The apparent dependence of these outcomes on particular physical parameters forms the basis of what physicists and philosophers call the

fine-tuning problem (Barnes, 2012; Lewis et al., 2026). The observation raises a further question: why do the parameters have the values they do?

Several possibilities have been proposed. Perhaps countless physical domains exist, with observers necessarily appearing only in the regions compatible with their existence. Perhaps the universe's life-permitting structure is better understood as the result of intentional selection. Perhaps a deeper physical theory will eventually derive the values from more fundamental principles, showing that what currently appears contingent could not have been otherwise. Another possibility is that the constants are simply facts about the universe for which no deeper explanation is required.

These proposals are often grouped together under the broad question of fine-tuning, but they do not perform the same explanatory task. A multiverse requires a physical account of the larger ensemble in which our universe is embedded. A teleological explanation introduces intentionality as part of the explanation. A deeper physical theory would need to derive the observed values rather than merely reproduce them. A brute-fact account questions whether an additional explanation is necessary at all.

The fine-tuning problem therefore begins with a physical question that can be investigated directly: how sensitive is the emergence of complexity to the universe's fundamental parameters? The more difficult question concerns interpretation. Physics can tell us what changes when the parameters change. It cannot, from that fact alone, determine why the parameters possess their observed values.

The Physics of a Life-Permitting Universe

Fine-tuning is sometimes presented as though it were a single numerical measurement expressing the improbability of the universe. The scientific literature describes something broader. Fine-tuning arguments examine the ways in which complex structures depend upon physical laws, constants, and initial conditions.

Researchers can vary parameters within theoretical models and investigate the consequences. If a modest change prevents a stable structure from forming or significantly alters cosmic evolution, the relevant parameter is sensitive to that outcome. The result is scientifically meaningful even before any explanation for the observed value is proposed. It describes the behavior of a physical model.

Barnes's (2012) extensive review examines fine-tuning arguments involving the cosmological constant, galaxy formation, stellar structure, particle physics, chemistry, entropy, inflation, and the dimensionality of space and time. Lewis et al. (2026) revisit these questions from a more contemporary perspective while examining proposed explanations such as theism and the multiverse. Taken together, the literature presents fine-tuning as a broad pattern of dependence: the physical parameters of the universe constrain what kinds of structures can emerge and persist. Several examples make that dependence clearer.

The Cosmological Constant

The cosmological constant, represented by Λ, is one of the clearest examples of fine-tuning because it affects the expansion of the universe on very large scales. A positive cosmological constant contributes to the accelerated expansion observed today. Its measured value is also extraordinarily small when compared with some theoretical expectations arising from quantum field theory, a discrepancy associated with the cosmological constant problem.

Fine-tuning raises a somewhat different question: what would happen if Λ were substantially larger?

Weinberg's (1987) analysis showed that a sufficiently large positive cosmological constant would prevent matter from becoming gravitationally bound into structures such as galaxies. If cosmic expansion became dominant too early, gravity would have insufficient time to gather matter into the structures from which stars and planets eventually form. The result establishes a physical constraint on the range of Λ compatible with large-scale structure.

The significance of Weinberg's argument lies partly in what it does **not** establish. It does not demonstrate that the cosmological constant was designed, nor does it establish that multiple universes exist. It shows that structure formation imposes an anthropic bound on Λ. The physical calculation therefore identifies a narrow region of parameter space relevant to complexity without, by itself, explaining why the observed value lies within that region.

Nuclear Physics and the Origin of Carbon

The dependence of complex chemistry on fundamental physics becomes especially apparent inside stars. Nuclear reactions determine which elements can be synthesized, and those elements later become the raw materials for chemical processes in planets and living systems. Carbon is particularly important because of its central role in terrestrial biochemistry.

Hoyle's (1954) work on stellar nucleosynthesis examined the nuclear reactions occurring in very hot stars and their role in producing elements from carbon through nickel. Later work in the fine-tuning literature has emphasized the sensitivity of carbon and oxygen production to the relationships among nuclear parameters. One well-known example is the excited state of carbon-12 known as the Hoyle state, which plays an important role in the triple-alpha process through which stars produce carbon (Barnes, 2012).

The discussion is sometimes reduced to a claim that a specific percentage change in the strength of the strong nuclear force would eliminate carbon. Such statements can obscure the underlying physics. The outcome depends upon which parameters are varied, which others are held fixed, and what counts as a viable alternative chemistry. A more careful conclusion is that complex chemistry depends upon nuclear processes, while those processes depend upon relationships among fundamental physical quantities. Alter those relationships, and the pathways through which stars manufacture the elements needed for complex chemistry can change substantially.

That finding contributes to the fine-tuning problem because it extends the dependence across several stages of cosmic development. The possibility of complex chemistry is tied to the behavior of stars, the behavior of stars is tied to nuclear physics, and nuclear physics is tied to

relationships among more fundamental parameters. The chain is physically meaningful even though it does not, by itself, provide a probability that the observed parameters were unlikely.

Initial Conditions and Cosmic Structure

Fine-tuning arguments also concern the initial state of the universe. The early universe appears to have begun in an extraordinarily special low-entropy state, producing what has become known as the problem of the "past hypothesis" (Barnes, 2012).

This question differs from the usual discussion of fundamental constants. A constant forms part of the laws used to describe a physical system, whereas an initial condition specifies the state in which that system begins. The distinction matters because a theory that explains one may leave the other untouched.

If several apparently special features characterize the universe, a satisfactory explanation would need to determine how those features are related. They could be independent conditions. They could emerge from a deeper theory. They could result from a broader physical process. They could simply be facts about the universe we inhabit. The existence of several fine-tuning problems therefore does not guarantee that a single explanation will account for all of them.

The Anthropic Principle: Selection Without a Complete Explanation

The anthropic principle enters the debate because observations are necessarily conditioned by the existence of observers. Carter's (1974) formulation emphasized that observers can only find themselves in environments compatible with their existence. At its most basic level, the principle is almost unavoidable: a universe incapable of producing observers cannot contain anyone asking questions about it.

The principle becomes more significant when applied to cosmological parameters. Suppose a parameter strongly affects galaxy formation. An observer who exists in a galaxy already has information about that parameter's value because the observer's existence requires the relevant structures to have formed. In this sense, observation is subject to a selection effect.

Bostrom's (2002) work on anthropic reasoning develops this issue through observation-selection effects. Evidence available to an observer may depend upon the fact that the observer exists in the first place. Such selection effects can alter how evidence should be interpreted, particularly in hypotheses involving many observers distributed across different environments. The anthropic principle, however, does not by itself provide a complete causal explanation. Knowing that observers can exist only under certain conditions tells us why we should expect to observe compatible conditions. It does not establish why a universe capable of producing observers exists at all. Selection changes what an observer should expect to see; it does not, without additional premises, explain the origin of the conditions being observed.

When Selection Becomes a Multiverse

The anthropic argument becomes more substantial when combined with the possibility that many physical environments exist. Multiverse proposals suggest that the observable universe

may constitute one region within a much larger physical reality. Some versions arise from inflationary cosmology, in which inflation can continue in some regions while ending in others, potentially producing a vast collection of expanding regions with different physical properties.

The details are crucial. There is no single multiverse theory, and inflation does not automatically entail every version of a multiverse. Different inflationary models produce different kinds of domains and different possibilities for variation in physical parameters. Garriga and Vilenkin (2001), for example, examined inflationary models involving an enormous ensemble of regions and explored their statistical implications.

The attraction of this approach is straightforward. If many physical domains exist and only some permit complex life, observers should naturally expect to find themselves in one of those domains. A condition that looks extraordinary in a single universe may be less surprising when considered within a sufficiently large ensemble.

That reasoning depends upon a substantial assumption, however: the ensemble must possess a meaningful probability structure. The existence of many universes does not itself tell us how frequently different kinds of universes occur. Researchers must specify which domains belong to the ensemble and how they should be weighted. In infinite or extremely large ensembles, different choices of measure can lead to different probabilities, making the measure problem one of the central difficulties for anthropic multiverse reasoning (Lewis et al., 2026).

The multiverse therefore offers a physical setting in which selection effects may become explanatory, but it does not remove every unanswered question. Why does the larger ensemble exist? What determines its possible range of physical conditions? Which probability measure should be used? The proposal shifts the explanatory problem outward. That may be a legitimate move, but it is not the same as solving every part of the original problem.

When Fine-Tuning Becomes an Argument About Purpose

A teleological response takes the same physical evidence in another direction. Rather than asking whether additional universes could make a life-permitting universe statistically expected, it asks whether the observed order is better explained by intentional selection.

Collins (2009) develops one of the most extensive versions of this argument. His approach treats fine-tuning as evidence that can support theism through comparative explanatory reasoning. If a universe capable of supporting complex life is more expected under a hypothesis involving purposive intelligence than under certain alternatives, then the observation may count as evidence for that hypothesis.

The argument does not depend upon identifying a designer through a physical measurement. Its force comes from an inference about explanation. Physical observations form the evidence, while the conclusion concerns what kind of explanation best accounts for those observations. That makes the argument philosophical rather than a scientific theory in the ordinary sense.

One major objection is that the proposed explanation may simply relocate the problem. If the universe requires explanation because its parameters are highly specific, why should an intelligent designer not require explanation as well?

A teleological defender can answer that the issue is comparative rather than absolute. The question is not whether a proposed explanation itself raises further questions, but whether it provides a better account of the evidence than its alternatives. On this view, considerations of simplicity, explanatory scope, and probability remain relevant even when the proposed explanation is metaphysical rather than physical (Collins, 2009).

Teleology and the multiverse therefore differ in kind. The multiverse extends the physical picture by proposing additional physical domains. Teleology extends the explanatory picture by appealing to intentionality. The two can be compared as explanations of the same evidence, but they should not be treated as rival physical theories.

The Possibility That Nothing More Is Required

There is also a simpler possibility: perhaps the constants simply have the values they have. A brute-fact approach does not require the claim that the universe is random. It makes a more limited claim: not every fact about reality necessarily has a deeper explanation. On this view, the observed values can serve as starting points for physical inquiry without demanding an ultimate reason for why those values exist.

The position can feel unsatisfying because scientific investigation is often driven by the discovery of deeper causes. A quantity that currently appears arbitrary may eventually emerge from a more fundamental symmetry, a unified theory, or another physical principle. But until such a theory exists, treating the measured values as fundamental is scientifically coherent.

The remaining difficulty is philosophical. If the universe's parameters are contingent facts, the fine-tuning question may simply stop there. Whether that is an adequate stopping point depends upon a broader view of explanation. Scientific inquiry can continue searching for deeper causes, but it cannot guarantee that every fact has another fact beneath it that explains it.

Four Ways of Understanding the Same Universe

The major approaches can be compared more clearly by asking what each one considers to be in need of explanation.

Table 1. Comparison of Proposed Explanations for the Apparent Fine-Tuning of the Universe

Framework	Central Idea	What It Attempts to Explain	Remaining Difficulty
Physical Necessity / Deeper Theory	The observed parameters may eventually follow from more fundamental physical laws.	Why the parameters are not genuinely independent.	No complete theory currently derives all of the relevant parameters in this way.

Multiverse / Anthropic Selection	Many physical domains exist, and observers arise only in domains compatible with their existence.	Why observers should find themselves in a life-permitting environment.	Requires a physically motivated ensemble and a defensible probability measure.
Teleological Explanation	Life-permitting conditions provide evidence of purposive selection.	Why the physical order is especially conducive to complex life.	Depends on metaphysical premises and comparative judgments about explanation.
Brute Fact	The parameters simply possess their observed values.	Nothing beyond the physical fact itself.	Leaves the deeper question of why these values exist unanswered.

The approaches diverge at the level of the question they are trying to answer. A deeper physical theory asks why the parameters have to be what they are. The multiverse asks why observers should find themselves in a compatible environment if many environments exist. Teleology asks whether the order evident in the physical world points toward intention. The brute-fact position questions whether any further explanation should be expected.

Much of the confusion surrounding fine-tuning comes from moving between these questions without marking the transition. A physical constraint on a parameter is one kind of result. A probability claim is another. An explanation for probability is another still.

The Problem With Calling Fine-Tuning "Improbable"

The language of improbability appears throughout discussions of fine-tuning, but it requires careful use. Suppose a researcher determines that a parameter must lie within a narrow interval for stars to form. The calculation establishes a narrow life-permitting region within the model. It does not automatically establish that the observed value was unlikely.

To speak meaningfully about probability, we need additional information. What alternative values were possible? How were those values distributed? Could the parameter vary independently of the other constants? What physical theory defines the space of alternatives?

Without answers to those questions, dramatic numerical claims about the universe being "one in a trillion" or "one in 10^50" do not possess a clear statistical meaning. Fine-tuning calculations can demonstrate sensitivity without establishing an objective probability distribution over all possible universes (Barnes, 2012; Collins, 2009).

This does not make the calculations unimportant. The strongest version of the fine-tuning observation need not depend upon an exact numerical probability. It is enough to establish that complex structure appears confined to a particular region of physical parameter space and that many specified changes would produce radically different outcomes. That physical fact creates an explanatory question even when the probability of the observed values cannot be calculated.

Parameter Independence

Many fine-tuning calculations vary one constant while holding the others fixed. This is useful because it isolates a particular relationship, but it may not correspond to the possibilities permitted by a deeper physical theory. If several constants are mathematically related, changing one without changing the others might describe a universe that could never physically exist.

Fine-tuning calculations therefore demonstrate sensitivity under specified conditions. They do not necessarily map the full landscape of physically possible universes (Barnes, 2012).

The Definition of Life

Fine-tuning arguments generally use the requirements of complex chemistry and terrestrial life as the reference point. That is reasonable because carbon-based life is the only life currently observed. It remains an assumption, however.

We do not know whether every conceivable form of complexity would depend upon exactly the same physical conditions. The phrase "fine-tuned for life" therefore carries less certainty when it is expanded into a claim about every possible form of existence. The broader the conclusion, the more assumptions it requires.

The Problem of Initial Conditions

The laws of physics and the initial state of the universe present separate explanatory problems. A theory might derive today's constants while leaving the early universe's initial conditions unexplained. Conversely, a theory of initial conditions would not automatically explain why the fundamental laws take their particular form.

The desire for one comprehensive explanation is understandable, but nothing about the structure of the problem guarantees that nature provides one (Barnes, 2012).

The Difference Between Observation and Explanation

The most important limitation may be the simplest one. We can establish that stars depend upon particular physical relationships, that the cosmological constant affects structure formation, and that nuclear physics influences the production of elements needed for complex chemistry. From these observations, however, additional conclusions require additional arguments.

"The universe was designed," "there are infinitely many universes," and "nothing deeper explains the constants" are all interpretive claims. None follows automatically from the physical calculations themselves. The calculations provide the evidence from which such arguments begin; they do not settle the arguments in advance (Lewis et al., 2026).

A Boundary Rather Than a Verdict

Fine-tuning is sometimes presented as a conflict between science and religion, with cosmology expected to determine whether design or the multiverse ultimately wins. The actual literature is less tidy.

The multiverse is connected to developments in theoretical cosmology, particularly inflationary models, but specific scenarios depend upon additional assumptions and face major problems involving probability, measure, and testability. Teleological reasoning does not occupy the same methodological category as a physical theory, yet it can still be evaluated as a philosophical argument concerning evidence and explanation. A deeper physical theory could transform the problem by showing that parameters currently treated as independent are in fact necessary. A brute-fact account remains possible if no deeper explanation is required (Lewis et al., 2026).

These possibilities do not have identical epistemic status, and there is no reason they should. They begin with the same physical observation and then move in different directions. The multiverse extends the physical setting. Teleology extends the explanatory framework. Physical necessity searches beneath current laws. Brute fact declines to assume that a further explanation must exist.

That division helps clarify why fine-tuning occupies a boundary between disciplines. The underlying physics can be studied through models, calculations, and observations. The significance assigned to that physics depends partly on assumptions about probability, explanation, and what kinds of causes are admissible.

Conclusion: The Question That Remains

Fine-tuning provides a legitimate scientific reason to examine how strongly the possibility of complex structure depends upon the universe's physical parameters. The cosmological constant affects whether large-scale structure can form. Nuclear physics determines whether stars can manufacture elements needed for complex chemistry. The early universe's initial conditions influence the structures that emerge later. These relationships are physical and can be investigated independently of the philosophical conclusions drawn from them.

The uncertainty begins when those relationships are interpreted. A life-permitting universe is not automatically an improbable universe. Establishing improbability requires a defined set of alternatives and some means of assigning probabilities to them, and no universally accepted probability distribution exists for the relevant space of possible universes. The physical sensitivity of the parameters is therefore easier to establish than the probability of obtaining their observed values.

The same distinction applies to proposed explanations. A multiverse could make anthropic selection meaningful, but it would require a physically motivated account of the larger ensemble and a defensible way to assign probabilities within it. A teleological explanation could interpret the apparent order of the universe as evidence of intentionality, but that conclusion depends upon philosophical premises that extend beyond the physical measurements themselves.

A deeper physical theory could show that the apparent freedom of the constants is temporary, while a brute-fact position could maintain that no deeper explanation is necessary.

The persistence of the fine-tuning debate follows from this gap. Current physics allows us to study what the universe would look like under different physical conditions, yet it does not provide a settled answer to why the actual conditions take the values they do. That leaves several possibilities open without making them equally established.

Cosmology can determine the physical conditions under which galaxies, stars, chemistry, and complexity emerge. Anthropic reasoning can examine how the existence of observers affects what can be observed. Multiverse theories can ask whether our universe forms part of a broader physical ensemble. Teleological arguments can ask whether the observed order is better understood as purposive. A future physical theory might connect these questions in ways that are currently unavailable, or the deeper explanation may turn out to be less complete than we expect.

For now, the most defensible conclusion is also the most restrained. The universe appears to occupy a region of physical parameter space unusually compatible with complex structure, and changes to several fundamental parameters would produce dramatically different outcomes. That fact is scientifically significant. What it ultimately means remains an open question at the boundary of cosmology and philosophy.

Works Cited

Barnes, L. A. (2012). The fine-tuning of the universe for intelligent life. Publications of the Astronomical Society of Australia, 29(4), 529–564. https://doi.org/10.1071/AS12015

Bostrom, N. (2002). Anthropic bias: Observation selection effects in science and philosophy. Routledge.

Carter, B. (1974). Large number coincidences and the anthropic principle in cosmology. In M. S. Longair (Ed.), Confrontation of cosmological theories with observational data (IAU Symposium No. 63, pp. 291–298). Reidel.

Collins, R. (2009). The teleological argument: An exploration of the fine-tuning of the universe. In W. L. Craig & J. P. Moreland (Eds.), The Blackwell companion to natural theology (pp. 202–281). Wiley-Blackwell. https://doi.org/10.1002/9781444308334.ch4

Garriga, J., & Vilenkin, A. (2001). Many worlds in one. Physical Review D, 64(4), 043511. https://doi.org/10.1103/PhysRevD.64.043511

Hoyle, F. (1954). On nuclear reactions occurring in very hot stars. I. The synthesis of elements from carbon to nickel. The Astrophysical Journal Supplement Series, 1, 121–146. https://doi.org/10.1086/190005

Lewis, G., Barnes, L. A., & Goff, P. (2026). Cosmological fine-tuning: The view from 2025. Religious Studies, 62(S1), S100–S110. https://doi.org/10.1017/S0034412525101170

Weinberg, S. (1987). Anthropic bound on the cosmological constant. Physical Review Letters, 59(22), 2607–2610. https://doi.org/10.1103/PhysRevLett.59.2607

A Review of Sustainable Battery Design: Alternatives to Lithium-Ion Batteries: Case Studies of Redox, Na^+, Mg^{2+} Ion Based Battery

By Shloke Kusoorkar, Lawrence T. Tom

Abstract

The growing demand for rechargeable energy storage has increased interest in battery chemistries that may reduce the material, environmental and supply-chain pressures associated with lithium-ion systems. This study compares sodium-ion, magnesium-based and redox-flow batteries with lithium-ion technology using evidence drawn from published studies and technical reports. The evidence was organised into eight criteria: material abundance, environmental sustainability, safety, cycle life, commercial readiness, operating-temperature performance, normalised cost performance and gravimetric energy density. These criteria were evaluated using a sustainability-led multi-criteria decision analysis (MCDA), with environmental sustainability assigned a weight of 0.35 and material abundance a weight of 0.20. Alternative weighting scenarios and application-specific analyses were used to test the sensitivity of the rankings. Under the sustainability-led baseline, magnesium-based batteries achieved the highest composite score (8.15), followed by sodium-ion (7.37), redox-flow (7.12) and lithium-ion batteries (2.54). Increasing the environmental-sustainability weight to 0.40 strengthened the relative position of sodium-ion batteries, whereas balanced and reliability-focused scenarios placed redox-flow batteries ahead of sodium-ion. Application-specific MCDA identified redox flow as the strongest option for grid storage. Magnesium-based systems ranked highest in the modelled portable-electronics, electric-vehicle and low-cost-storage scenarios. However, these results were influenced by research-stage data and did not fully represent practical constraints related to packaging, balance-of-system requirements and manufacturing maturity. Within the scope of the literature and criteria examined, no single alternative battery technology emerged as a universal replacement for lithium-ion batteries across portable electronics, electric vehicles and grid storage. Lithium-ion batteries remain the most established option for energy-dense mobile applications. Redox-flow batteries appear better suited to long-duration stationary storage, while sodium-ion batteries offer a credible route to material and supply-chain diversification. Magnesium-based batteries remain a promising but comparatively immature research pathway.

Introduction

Rechargeable batteries are increasingly central to transport, communications and electricity systems. This study began with a practical design question: can sodium-ion, magnesium-ion and redox-flow technologies provide more sustainable alternatives to lithium-ion batteries, and which alternative is most suitable for electric vehicles, portable electronics such as laptops and mobile phones, and grid storage? Although the original question referred to alternative battery components, the comparison is conducted at the battery-technology level because sustainability and application suitability emerge from the interaction of electrodes, electrolytes, separators, cell architecture, manufacturing and balance-of-system requirements.

Lithium-ion batteries are used as the benchmark because they currently combine high full-cell energy density, high efficiency and mature manufacturing at commercial scale.

The rapid expansion of battery production has sharpened questions about mineral extraction, processing energy, geographical concentration and waste management (Whittingham; *Why Are Lithium-Ion Batteries, and Not Some Other Kind of Battery, Used in Electric Cars and Grid-Scale Energy Storage?*). Spent lithium-ion batteries can release metals, fluorinated compounds and organic electrolytes if poorly handled, and the environmental burden of primary lithium depends strongly on the deposit, extraction route and electricity mix ("Batteries - Why Lithium-Ion?"; Agusdinata et al.; Kelly et al.; Vera et al.). Regulation is consequently shifting from a narrow focus on in-use emissions towards carbon footprints, due diligence, recycled content and end-of-life recovery (*Regulation - 2023/1542 - EN - Batteries Regulation - EUR-Lex*). These concerns do not imply that lithium is geologically exhausted or that all lithium-ion systems are unsustainable; rather, they motivate a broader comparison of alternative chemistries across the full lifecycle.

Alternative battery technologies and application requirements

Sodium-ion batteries use a rocking-chair mechanism broadly analogous to lithium-ion systems but replace lithium with the larger and heavier sodium ion. Sodium is widely distributed and some sodium-ion designs avoid nickel, cobalt and copper, potentially reducing exposure to concentrated supply chains. Their generally lower energy density makes them less attractive for mass- and volume-constrained electronics, but they may be suitable for lower-cost mobility, stationary storage and applications in which resource availability and supply-chain diversification are prioritised (Usiskin et al.).

Magnesium-based batteries are attractive because magnesium is abundant and a metal electrode can, in principle, transfer two electrons (Usiskin et al.) per atom. Yet divalency also increases interactions with solvents, anions and host lattices (Manthiram; Matsui et al.). Slow solid-state transport, difficult desolvation, interfacial passivation and electrolyte instability remain central barriers. These limitations are particularly important for electric vehicles and portable electronics, where high power, fast charging, long cycle life and practical full-cell energy density must be demonstrated simultaneously (Medina et al.; Mohtadi and Mizuno). The term magnesium-based is used here because the evidence base may include magnesium-ion intercalation, magnesium-metal and magnesium-sulfur systems; these architectures should be separated in future chemistry-specific analyses (Nolte et al.; Olabi et al.).

Redox-flow batteries store electroactive species in externally circulated electrolytes. Their architecture separates, at least partly, the scaling of power in the cell stack from energy capacity in the tanks (Song et al.; Alotto et al.). This supports long-duration stationary and grid storage, long cycle life and potentially favourable safety performance, but requires pumps, tanks, piping and control systems. These balance-of-plant requirements make conventional redox-flow systems unsuitable for most laptops, mobile phones and passenger electric vehicles, even when an aggregate score appears competitive (Darling et al.; Skyllas-Kazacos et al.).

Research gap and study approach

There is growing interest in sodium-ion, magnesium-based and redox-flow batteries as possible alternatives to lithium-ion batteries. However, most existing studies examine these technologies separately or focus on only one performance measure, such as energy density, cycle life, cost or material abundance. This makes it difficult to judge which alternative is genuinely more sustainable when both environmental considerations and practical application requirements are taken into account.

Another challenge is that published studies often report results at different levels. Some values are theoretical or based on active materials, while others come from laboratory cells, commercial cells or complete battery systems. These values cannot always be compared directly. In addition, a battery that appears suitable for grid storage may not be appropriate for electric vehicles, laptops or mobile phones, where mass, volume and energy density are much more important.

This study addresses these gaps by comparing lithium-ion, sodium-ion, magnesium-based and redox-flow batteries within one common framework. The technologies are evaluated using eight criteria covering sustainability, technical performance, cost and commercial readiness. The comparison is then repeated for different applications, including portable electronics, electric vehicles, grid storage and low-cost stationary storage. Sensitivity and leave-one-out analyses are also used to test whether the rankings depend too strongly on particular weights or individual criteria. The aim is not to identify one universal replacement for lithium-ion batteries, but to determine where each alternative may offer a meaningful advantage and where further evidence is still needed.

Results

Literature-derived evidence

All data used in this study was obtained from published papers, technical reports, and where commercial status was relevant, manufacturer or industry sources. No experimental battery testing was carried out. Collecting the data was challenging because different studies reported performance in different ways. Some papers presented theoretical material values, while others reported results from laboratory cells, commercial cells or complete battery systems. The units, test temperatures, cycling conditions and definitions of cost also varied between sources. In order to make the comparison possible, the available values were grouped under eight criteria and converted to a common scale from 1 to 10. A score of 10 represented the most favourable value within the technologies included in this study, while a score of 1 represented the least favourable value. These scores are therefore relative values based on the collected literature and should not be treated as direct measurements of battery performance.

Performance profiles across criteria

The four battery technologies showed different strengths and weaknesses across the selected criteria (Fig. 1 and Supplementary Fig. 1). Lithium-ion batteries had the highest scores for energy density and commercial readiness, both with a score of 10. These are important advantages for applications such as electric vehicles, laptops and mobile phones. However, lithium-ion received low scores for material abundance, sustainability, safety and cycle life in the present dataset. Magnesium-based batteries received high scores for material abundance, operating-temperature performance and normalised cost performance. They also scored relatively well for sustainability, safety, commercial readiness and energy density. Their weakest result was cycle life, with a score of 1.8. Sodium-ion batteries performed strongly in material abundance and sustainability, both with scores of 10. They also had a high safety score of 8.5. Their lower scores were found in cycle life, cost performance, commercial readiness and operating-temperature performance. Redox-flow batteries achieved the highest scores for safety and cycle life. They also performed well for sustainability and commercial readiness. Their main weaknesses were low energy density and low normalised cost performance.

These results show that no technology performed best in every category (Supplementary Data Fig. 2b). The result already suggests an application-dependent outcome: compact mobile systems favour energy density and maturity, whereas stationary storage can accept lower energy density in exchange for safety, lifetime and storage duration.

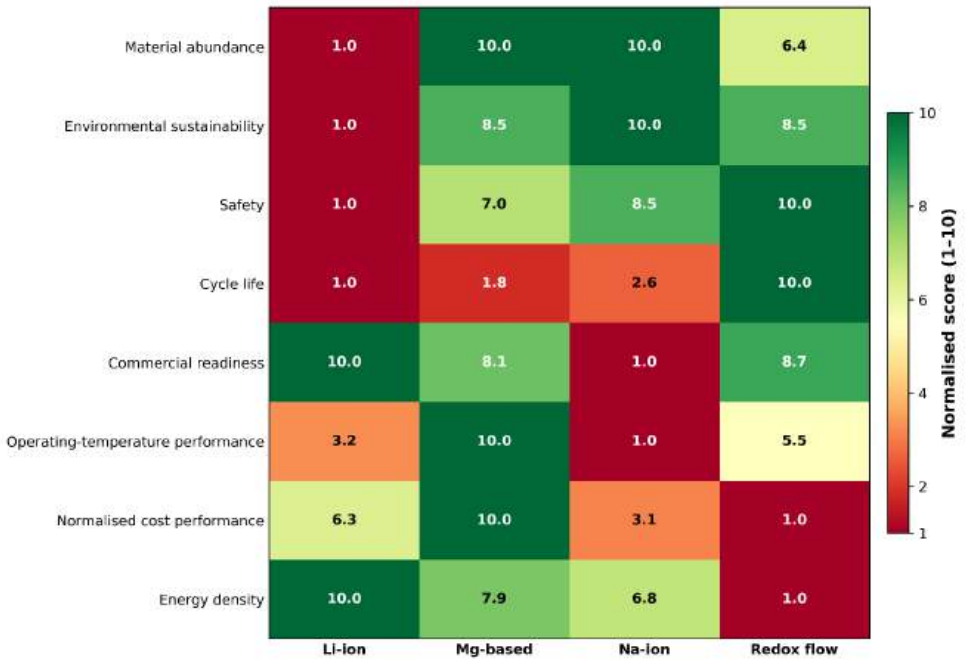


Figure 1: Normalized performance profiles across evaluation criteria. Heatmap of relative 1–10 scores for lithium-ion, magnesium-based, sodium-ion and redox-flow batteries. Green indicates comparatively favourable performance and red indicates comparatively unfavourable performance within the analysed dataset. Scores are dimensionless and should not be interpreted as proportional physical differences.

The unweighted average scores were 7.91 for magnesium-based batteries, 6.39 for redox-flow batteries, 5.38 for sodium-ion batteries and 4.19 for lithium-ion batteries (Supplementary Data Fig. 2a). Magnesium-based batteries also had the lowest standard deviation (2.71), which suggests that their scores were more evenly spread across the eight criteria (Supplementary Data Fig. 3a). However, the unweighted average gives the same importance to every criterion. This does not reflect the fact that energy density may be very important for a phone or electric vehicle but much less important for a stationary grid-storage system.

Effect of different weighting scenarios

The main analysis was structured around the sustainability focus of the study. Environmental sustainability was assigned a weight of 0.35 and material abundance a weight of 0.20, while the remaining weight was distributed across safety, cycle life, normalised cost performance, commercial readiness, operating-temperature performance and energy density. Under this weighting, magnesium-based batteries achieved the highest composite score of 8.15, followed by sodium-ion batteries at 7.37, redox-flow batteries at 7.12 and lithium-ion batteries at 2.54 (Fig. 2 and Table 1).

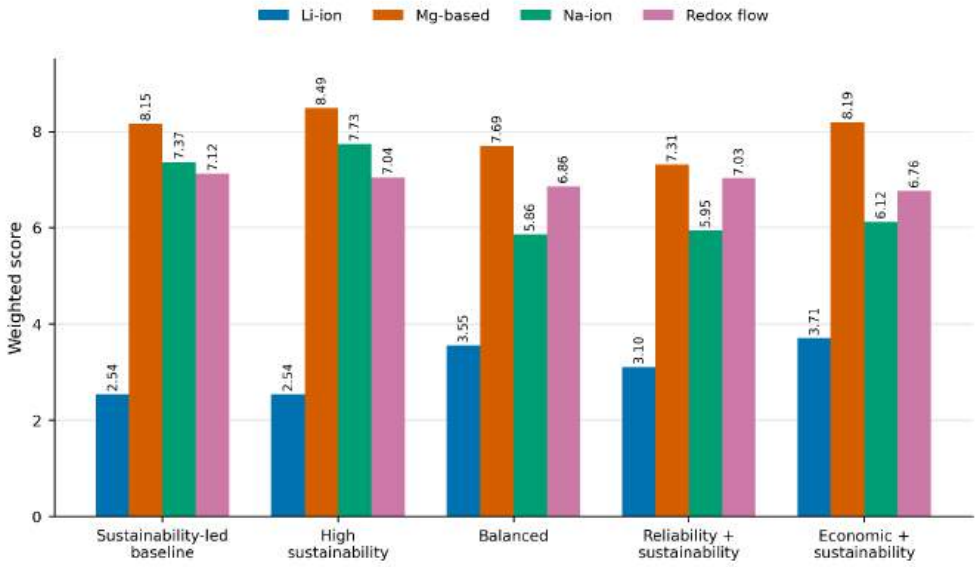


Figure 2**:** Sensitivity to general weighting choices. Weighted scores under the sustainability-led baseline, a higher-sustainability case, the original balanced case, a reliability-and-sustainability case, and an economic-and-sustainability case. Magnesium-based batteries remain first, whereas sodium ion and redox flow exchange second and third place as the emphasis changes.

The sensitivity analysis showed that the relative positions of sodium-ion and redox-flow batteries were strongly influenced by the weighting assumptions. Increasing the environmental-sustainability weight to 0.40 raised the magnesium-based score to 8.49 and the sodium-ion score to 7.73. By contrast, redox-flow batteries ranked above sodium-ion under the balanced, reliability-and-sustainability, and economic-and-sustainability scenarios. Magnesium-based batteries remained first and lithium-ion batteries remained fourth across all

general scenarios. The main variation therefore occurred between sodium-ion and redox-flow technologies: sodium-ion performed more strongly when environmental sustainability and material abundance were prioritised, whereas redox flow benefited when greater emphasis was placed on safety, cycle life and reliability.

Technology	Sustainability-led baseline	High sustainability	Balanced	Reliability + sustainability	Economic + sustainability
Lithium ion	2.54	2.54	3.55	3.10	3.71
Magnesium based	8.15	8.49	7.69	7.31	8.19
Sodium ion	7.37	7.73	5.86	5.95	6.12
Redox flow	7.12	7.04	6.86	7.03	6.76

Table 1: Weighted scores under general preference scenarios

Values are weighted sums of the normalised criterion scores. Higher values indicate a more favourable result under the stated weighting assumptions.

Application-specific comparison

Application-specific weighting produced a clearer separation between the technologies (Fig. 3 and Table 2). Redox-flow batteries achieved the highest score for grid storage at 7.53, slightly above magnesium-based batteries at 7.24. This result is consistent with the main design requirements of stationary storage, where safety, cycle life and storage duration are generally more important than low mass or compact volume.

Magnesium-based batteries ranked first in the portable-electronics, electric-vehicle and low-cost-storage scenarios. These results require careful interpretation. Redox-flow batteries ranked second for both portable electronics and electric vehicles, despite the fact that conventional systems require tanks, pumps and additional balance-of-plant components that are not practical for these applications. This inconsistency highlights a limitation of the weighted model: a common set of numerical criteria cannot fully represent architecture-specific design constraints (Lyon; Schirn; Johnson; Pinto-Bautista et al.; Leung; Greenwood et al.). The strong magnesium-based result was also influenced by favourable literature-derived scores for cost, commercial readiness, operating-temperature performance and energy density. It should therefore be regarded as an indication of research potential rather than evidence of current commercial superiority (Mousavian and Bautin; Malachi et al.; Leong et al.; Orikasa et al., “High Energy Density Rechargeable Magnesium Battery Using Earth-Abundant and Non-Toxic Elements”; Aurbach et al.; Gao et al.; Ponrouch et al.; Bella et al.; Rasheev et al.).

Sodium-ion batteries ranked second for low-cost storage and performed more strongly when greater emphasis was placed on sustainability and material abundance. Their lower energy density remains a disadvantage for compact electronics and long-range electric vehicles, but their

material availability and supply-chain profile make them a credible option for applications where greater battery mass or volume can be tolerated (Schirber; Glushenkov; Paul et al.; Machín and Márquez).

For low-cost stationary storage, magnesium-based batteries had the highest score of 8.27, while sodium-ion batteries ranked second with 6.52. This result was mainly influenced by the strong abundance, sustainability and cost-performance scores assigned to magnesium. Since commercial magnesium battery costs are not yet well established, the result represents potential rather than confirmed commercial performance.

Technology	Portable electronics	Electric vehicles	Grid storage	Low-cost storage
Lithium ion	5.69	4.79	2.99	3.52
Magnesium based	7.33	7.05	7.24	8.27
Sodium ion	4.94	5.18	6.03	6.52
Redox flow	6.11	6.55	7.53	5.99

Table 2 | Application-specific weighted scores.

The application scenarios use the same normalised score matrix but different criterion weights. They do not include every system-level design constraint.

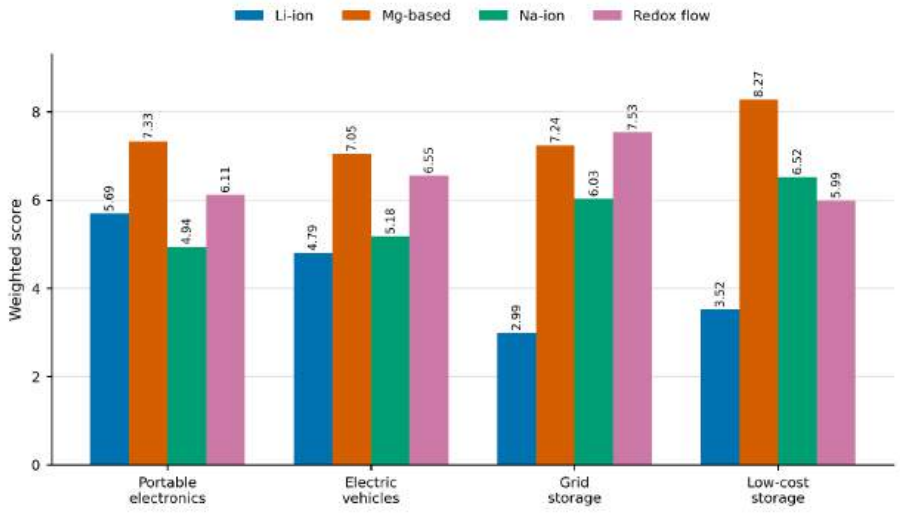


Figure 3: Application-specific weighted comparison. Composite scores for portable electronics, electric vehicles, grid storage and low-cost stationary storage. These are conditional model outputs based on the stated weights and should not be read as direct evidence of present commercial suitability.

Robustness and trade-offs

The contribution analysis shows how the weighted criteria shaped the baseline scores (Fig. 4a). Environmental sustainability made the largest contribution to the sodium-ion score, adding 3.50 points, while contributing 2.98 points to both magnesium-based and redox-flow batteries. Material abundance added a further 2.00 points to the magnesium-based and

sodium-ion totals. By comparison, lithium-ion batteries gained only 0.35 points from environmental sustainability and 0.20 points from material abundance. Since these were the two most heavily weighted criteria, their low contribution had a major effect on the lithium-ion baseline score. Redox-flow batteries gained more of their total from safety and cycle life, whereas the magnesium-based score was supported by relatively strong contributions across several criteria.

The leave-one-out analysis tested whether the ranking depended on any single criterion (Fig. 4b). Magnesium-based batteries remained the highest-ranked technology when each criterion was removed in turn. The narrowest separation occurred when normalised cost performance was excluded, with scores of 7.95 for magnesium-based, 7.84 for sodium-ion and 7.80 for redox-flow batteries. Removing cycle life increased the magnesium-based score to 8.86, indicating that cycle life was its largest disadvantage in the dataset.

Overall, the first-place ranking was stable within the weighting model, although the small score differences in several tests suggest that the result could still change if the underlying literature values or scoring assumptions were revised.

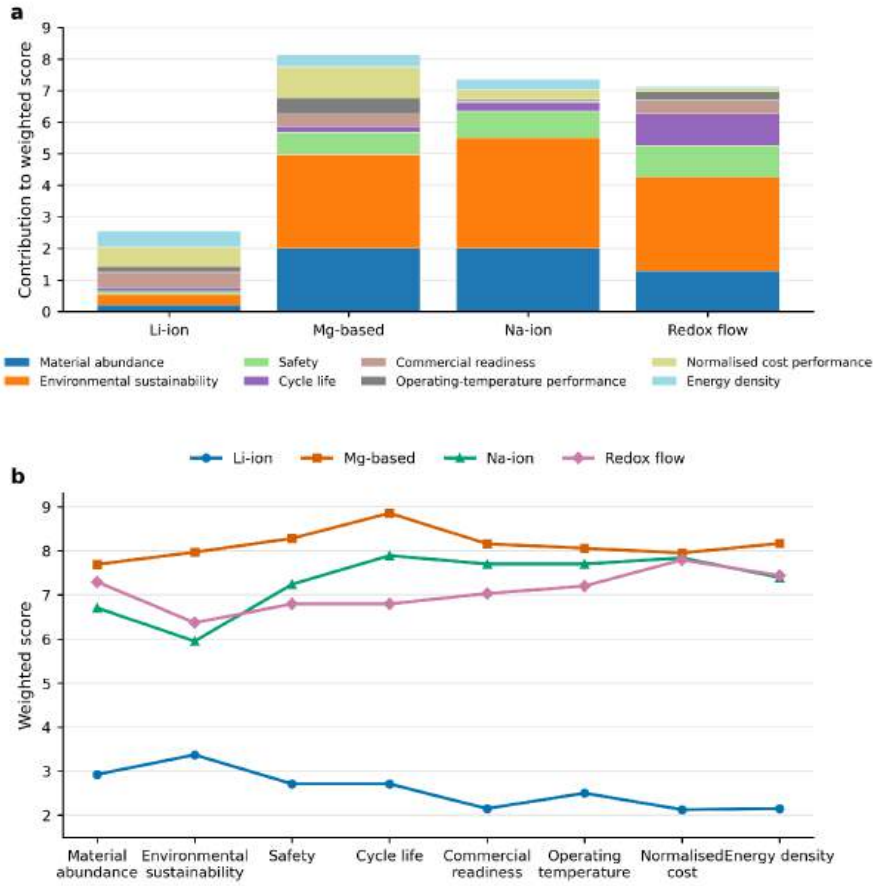


Figure 4: Contributions and robustness under the sustainability-led baseline. **a**, Contribution of each criterion to the weighted total. **b**, Scores after removing one criterion and renormalising the remaining weights. Magnesium-based batteries remain first in all eight tests, although the margin narrows when cost performance is excluded.

Main trade-offs between technologies

The pairwise plots show why one overall score is not enough (Fig. 5). Lithium ion combines high energy density and readiness with a low sustainability score in the present dataset. Redox flow shows the opposite pattern: strong sustainability, safety and cycle life but poor compactness. Sodium-ion sits between these extremes, with high sustainability and moderate energy density. Magnesium based appears strong on several pairs, although that position combines evidence from different magnesium architectures and development stages.
The readiness-versus-abundance comparison also exposes the least secure part of the dataset. Magnesium based appears both abundant and relatively mature, while sodium ion appears abundant but minimally ready. Those positions are difficult to reconcile with the different kinds and dates of evidence used. A common readiness rubric and a source-by-source audit are therefore needed before the absolute rankings are treated as final.

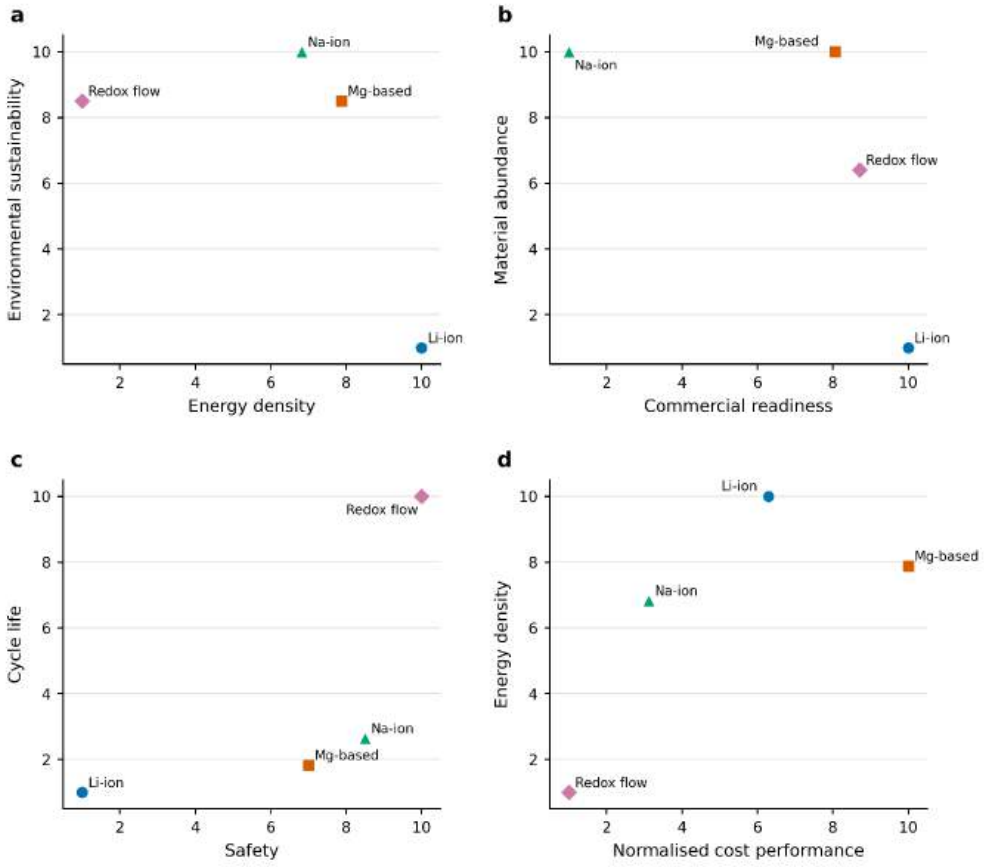


Figure 5: Pairwise trade-offs among normalized criteria. **a**, Energy density versus environmental sustainability. **b**, Commercial readiness versus material abundance. **c**, Safety versus cycle life. **d**, Normalized cost performance versus energy density. All axes show dimensionless scores. A higher cost-performance score denotes more favourable cost rather than a higher raw cost per kilowatt-hour.

Summary of the results

The analysis showed that the relative value of each battery technology depends on the intended application and on the criteria given greatest importance. Redox-flow batteries provided

the strongest and most technically credible result for grid storage, mainly because of their favourable safety and cycle-life performance. Sodium-ion batteries were more competitive in scenarios that emphasised material availability, environmental sustainability and lower-cost design, suggesting a possible role in stationary storage and other applications where moderate energy density can be accepted.

Magnesium-based batteries achieved the highest composite scores across the main weighting scenarios. However, this result was influenced by several literature-derived values from laboratory and early-stage studies. The ranking should therefore be interpreted as evidence of technological potential rather than demonstrated commercial advantage. Lithium-ion batteries received a lower score under the sustainability-led framework, but they remained the most mature and practical option for applications requiring high energy density, compact design and established manufacturing, including mobile phones, laptops and mainstream electric vehicles. Overall, the findings support an application-specific portfolio of battery technologies rather than the replacement of lithium-ion batteries by a single alternative (Supplementary Fig. 1).

Discussion

The sustainability-led analysis places greater emphasis on the environmental and resource implications of battery design. Under this framework, sodium-ion batteries move into second place because environmental sustainability and material abundance together account for more than half of the baseline score. Redox-flow batteries remain close behind and rank above sodium-ion systems when greater weight is given to reliability, safety and cycle life. Magnesium-based batteries remain the highest-ranked option across the general scenarios, but this result should be interpreted as an indication of potential within the available evidence rather than as evidence of present commercial readiness.

The findings also show that a composite sustainability score cannot replace engineering judgement. Redox-flow batteries illustrate this clearly. Their high safety and cycle-life scores make them well suited to grid storage, where long operating life and storage duration are more important than compactness (Fig 1). The same architecture is not realistic for mobile phones, laptops or conventional passenger vehicles because it requires tanks, pumps and additional balance-of-system components. Lithium-ion batteries present the opposite case. They receive a relatively low score under the sustainability-led framework, yet they still provide the strongest proven combination of energy density, power capability, compact packaging and manufacturing maturity for mobile applications.

The analysis therefore supports an application-specific interpretation rather than a universal ranking. A technology may perform well against sustainability criteria but remain unsuitable for a particular use because of limitations in size, weight, power delivery, operating complexity or commercial maturity.

Technology implications

For portable electronics, lithium-ion batteries remain the practical benchmark. Their high energy density and compact design are difficult to replace in devices where mass and volume are tightly constrained. Sodium-ion batteries may become relevant in lower-cost products or applications where resource security is given greater importance, but their lower energy density remains a disadvantage. Magnesium-based batteries would require stronger evidence from practical full cells, particularly for cycle life, charging rate and manufacturability, before their modelled performance could support a design recommendation. Conventional redox-flow systems are not suitable for portable electronics because of their balance-of-plant requirements.

The same considerations apply to electric vehicles. A meaningful comparison must include pack-level energy density, power density, charging performance, thermal behaviour, safety, cycle life and production scale. Lithium-ion technology currently satisfies this combination of requirements more completely than the alternatives considered. Sodium-ion batteries may be suitable for shorter-range or lower-cost vehicles, where reduced energy density can be accepted in exchange for potential material and cost advantages. Magnesium-based batteries remain a longer-term research option, while conventional redox-flow systems are poorly matched to passenger-vehicle mass and volume constraints.

Grid storage offers the clearest opportunity for broader technology diversification. Redox-flow batteries benefit from long cycle life, favourable safety characteristics and the ability to increase storage capacity by adding more electrolyte. Sodium-ion batteries are also credible for stationary applications, especially where material availability, manufacturing compatibility and supply-chain diversification are important. Magnesium-based systems may eventually contribute to stationary storage, but their present position depends on research-stage values that require further validation (Winsberg et al.; L. Li et al.; R. Li et al.; Tolmachev; “New Report Showcases How Innovation Can Fast Track Affordable Energy Storage”; Nichols).

Taken together, the results support a portfolio approach to battery development. Lithium-ion batteries remain appropriate for energy-dense mobile applications, redox-flow batteries are better matched to long-duration stationary storage, sodium-ion systems offer a plausible route to resource diversification, and magnesium-based batteries remain a promising but less mature research pathway.

Limitations and future work

The main limitation of this study is that the data came from many different published sources, and these sources did not always report battery performance in the same way. Some studies focused on theoretical properties of battery materials, while others tested laboratory cells, commercial batteries or complete storage systems. Cycle-life results were also measured under different temperatures, charging rates and end-of-life conditions. Cost estimates often referred to different years and different parts of the battery system. Min–max normalisation made it possible to compare the data on the same scale, but it could not remove all the uncertainty caused by these differences.

The battery groups used in the analysis were also quite broad. Lithium-ion batteries include several chemistries with different levels of safety, lifetime and use of critical materials. The magnesium-based category included magnesium-ion, magnesium-metal and magnesium–sulphur batteries, which are not all at the same stage of development. Redox-flow batteries also differ depending on the electrolyte, membrane and system design. Grouping each technology into one category made the analysis easier, but it also hid some important differences within each group.

Another limitation is the choice of weights used in the MCDA. The weights were selected to reflect the sustainability focus of the study, but they were not based on a formal survey of experts or stakeholders. This means that the final ranking depended partly on the priorities built into the model. Environmental sustainability was kept separate from material abundance, safety, cost and cycle life to reduce the risk of counting the same advantage more than once. The balanced and application-specific scenarios helped show how the rankings changed when different priorities were used.

A stronger future study should include a complete source-to-score table. This table should show the original reference, battery chemistry, value, unit, test conditions, measurement level and quality of evidence behind every score. Commercial readiness should also be judged using the same scoring system for all technologies, and cost values should be converted to the same currency year and system level. Future work should compare specific battery chemistries rather than only broad technology groups.

Uncertainty should also be included more directly in the analysis. Instead of using one fixed value for each score and weight, future work could use a range of possible values and test them through Monte Carlo simulation. This would show how often each battery ranks first and how much the results change under different assumptions.

The results should therefore be treated as an early comparison of sustainability and application trade-offs. They show where each technology appears strongest in this study, which assumptions have the greatest effect on the rankings and where more reliable evidence is still needed.

Methods

Study design and literature search

This was a literature-based comparative study. No experimental battery testing was carried out. The aim was to compare lithium-ion, sodium-ion, magnesium-based and redox-flow batteries in relation to sustainability and their suitability for portable electronics, electric vehicles, grid storage and low-cost stationary storage.

The work began by identifying published information on the four battery technologies. A scoping-review approach, guided by the preferred reporting items for systematic reviews and meta-analyses extension for scoping reviews (PRISMA-ScR), was used because the available evidence covered a wide range of study types, including material studies, laboratory cells,

commercial batteries, technical reports and industry publications. The literature review was then combined with MCDA so that the technologies could be compared using the same set of criteria (Tricco et al.; Phillips; Orikasa et al., “High Energy Density Rechargeable Magnesium Battery Using Earth-Abundant and Non-Toxic Elements”; Wu et al.; Greitemeier et al.; Aurbach et al.; *Advances in Fire Suppression and Battery Safety*).

The final database search was carried out on 25th June, 2026. The databases searched were [1-55], and the main search period covered publications from 2015 to 2026. Earlier studies were retained where they were important for explaining the origin, electrochemical mechanism or development of battery technology.

Search terms combined the names of the battery technologies with terms related to performance, cost and sustainability. A representative search expression was:

> (“lithium-ion battery” OR “sodium-ion battery” OR “magnesium-ion battery” OR “magnesium battery” OR “magnesium-sulfur battery” OR “redox flow battery”) AND (“energy density” OR “cycle life” OR “cost” OR “safety” OR “operating temperature” OR “commercial readiness” OR “material abundance” OR “life cycle assessment” OR “sustainability”).

The search returned **67** records. After duplicates and unsuitable sources were removed, **55** sources were included in the final evidence base.

Selection of sources

A source was considered suitable when it clearly identified the battery chemistry and provided numerical data, or enough information to assign a score using a stated rubric. The reported unit, measurement level and test conditions also had to be clear enough for the result to be interpreted.

Sources were excluded when the reported values could not be traced, when material-level capacity was presented as though it were full-cell energy density, or when values from different system boundaries were combined without explanation. Promotional claims that could not be supported by technical evidence were also excluded from the scientific performance assessment. Manufacturer reports and news articles were used mainly to understand commercial activity, pilot production and deployment status. They were not used as the only evidence for electrochemical performance, safety or lifecycle sustainability.

Data extraction and comparison

For each source, the evidence table recorded the battery family, specific chemistry, reported value, unit, publication year, evidence type and measurement level. Evidence type distinguished theoretical, estimated and measured values, while measurement level identified whether the result referred to a material, electrode, cell, module, pack or complete system.

Test conditions were also recorded where available, including temperature, current density or C-rate, depth of discharge, cycle-life endpoint and reported uncertainty. Measured full-cell and commercial-cell data were prioritised in the main comparison. Theoretical and active-material results were retained to provide context and indicate possible future performance, but they were not treated as equivalent to commercial system data.

Gravimetric energy density was expressed in watt-hours per kilogram where possible, while volumetric energy density was recorded separately. Specific capacity in milliampere-hours per gram was not treated as energy density unless the average operating voltage and an appropriate full-cell mass basis were also available.

Cycle life was defined, where the published information allowed, as the number of full-equivalent cycles completed before capacity declined to 80% of its initial value. Cost data was more difficult to compare because published values referred to different currencies, price years and system boundaries. A final economic comparison therefore requires all cost values to be converted to a common currency year and clearly defined boundary, such as cell, pack or installed system.

Evaluation criteria and weighting

Eight criteria were used to link sustainable design with practical performance. The baseline was a weighted-sum multi-criteria decision analysis. Environmental sustainability received a weight of 0.35 and material abundance 0.20 because the main research question concerned sustainability. Safety, cycle life and normalised cost performance each received 0.10, while commercial readiness, operating-temperature performance and energy density each received 0.05. These weights were fixed before the updated results were calculated. They represent a transparent analytical choice rather than a universal or empirically derived set of priorities..

Criterion	Operational definition	Direction	Weight
Environmental sustainability	Lifecycle impacts excluding abundance, safety and cost where these are scored separately	Higher is better	0.35
Material abundance	Availability and geographical distribution of principal active and critical materials	Higher is better	0.20
Safety	Thermal stability, flammability, gas generation, short-circuit and abuse response	Higher is better	0.10
Cycle life	Full-equivalent cycles to the defined end-of-life threshold	Higher is better	0.10
Normalised cost performance	Relative cost score after reversing the direction of raw cost	Higher is better	0.10
Commercial readiness	Maturity from laboratory demonstration to scaled deployment	Higher is better	0.05

Criterion	Operational definition	Direction	Weight
Operating-temperature performance	Validated operating range and performance outside room temperature	Higher is better	0.05
Energy density	Measured full-cell gravimetric energy density	Higher is better	0.05

Table 3: Evaluation criteria and baseline weights

Weights sum to 1.00. Environmental sustainability should be defined so that material abundance, safety, cost and cycle life are not counted twice.

Commercial readiness was interpreted on a common 1–10 scale: 1–2 for fundamental research, 3–4 for laboratory full-cell validation, 5–6 for prototype or pilot production, 7–8 for early commercial deployment and 9–10 for mature mass production. The score matrix should be audited against this rubric before submission. The overall workflow used to extract, normalise and evaluate the battery data is summarised in Fig. 6.

Multi-Criteria Decision Analysis Workflow

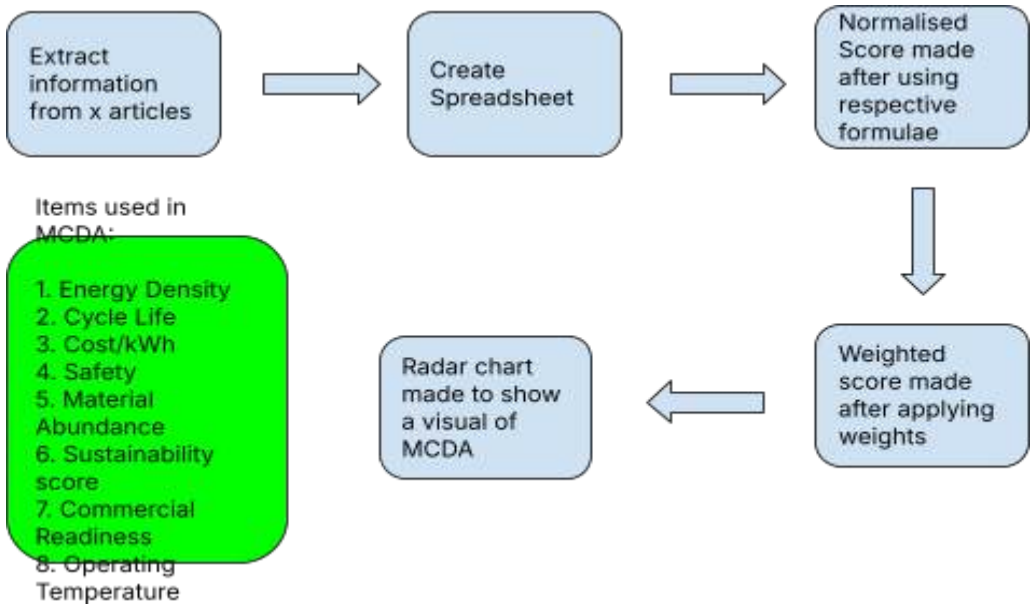


Figure 6: Workflow of the literature-based multi-criteria decision analysis. Relevant data were extracted from published studies and technical reports and compiled in a structured spreadsheet. The criterion values were converted to normalised scores using benefit- or cost-based min–max equations. Scenario-specific weights were then applied to calculate composite scores and compare battery technologies across general and application-specific conditions.

Normalisation and aggregation

The raw literature values used different units and ranges. Benefit criteria, where larger values were favourable, were converted to a 1–10 scale using:

$$S_i = 1 + 9[(x_i - x_{min})/(x_{max} - x_{min})] \quad (1)$$

For cost criteria, where a smaller raw value was favourable, the direction was reversed:

$$S_{i,j} = 1 + 9[(x_{j,max} - x_{i,j})/(x_{j,max} - x_{j,min})] \quad (2)$$

Here, $S_{i,j}$ is the normalised score for technology i under criterion j, $x_{i,j}$ is the raw value, and $x_{j,min}$ and $x_{j,max}$ are the lowest and highest values included for that criterion. A score of 10 therefore represents the most favourable value in this four-technology comparison and 1 the least favourable. If all technologies had the same value, the criterion was treated as non-discriminating.

The total weighted score was calculated as:

$$T_i = \sum_{j=1}^{m} w_j S_{i,j} \quad (3)$$

T_i is the total score, w_j is the criterion weight and m is the number of criteria. The method is compensatory, so a strong result in one criterion can offset a weak result in another. For that reason, totals were interpreted alongside individual scores, contribution plots and robustness tests.

Sensitivity analyses

Four alternative general scenarios were evaluated. The high-sustainability case increased environmental sustainability to 0.40. The balanced case reproduced the original, more even weights. Reliability-and-sustainability and economic-and-sustainability cases shifted emphasis towards lifetime and cost, respectively. Four application scenarios represented portable electronics, electric vehicles, grid storage and low-cost stationary storage. Their weights were chosen from the design priorities described in the literature: compact energy and maturity for electronics; energy, safety, life, cost and readiness for vehicles; and safety, life, sustainability and duration-related economics for grid storage.

A leave-one-out analysis removed each criterion and renormalised the remaining weights to sum to one. The analysis did not propagate uncertainty in the literature values. A confirmatory study should use value ranges and plausible weight distributions in a Monte Carlo analysis and report the probability of each technology ranking first, together with median ranks and uncertainty intervals.

Software and reproducibility

Data cleaning, weighted scoring, ranking, sensitivity analysis and figure generation were carried out in Python. The pandas and NumPy libraries were used for data handling and numerical calculations, while Matplotlib was used to produce the figures. To support transparency and reproducibility, the input score matrix, weighting schemes, output tables and analysis code should be deposited in a permanent public repository before submission. Cost-related figures are labelled “normalised cost performance” because they present dimensionless comparative scores rather than actual monetary costs per kilowatt-hour.

Code availability

The Python code used for data cleaning, normalization, sensitivity analysis and figure generation is attached analysis_script.

Acknowledgements

The content of this paper and my growth in the field of science and engineering has received priceless support from many persons.

First, my family. My parents have raised and given me the education and opportunity to be able to create and complete a formal paper such as this one.

Next, I give my greatest gratitude to my mentor, Dr. Lawrence T. Tom. Dr. Lawrence gave me the necessary and fullest guidance in order to make a proper paper. Through Dr. Lawrence's PhD experience of tactical skills and research architecture, I was able to have a greater appreciation and humility for the world of research. I am forever grateful for the inspiration that I gained through the experience he has provided me.

Lastly, I am grateful to Lumiere Education who provided the platform and guidance in order to provide the opportunity to create a paper like this.

Work Cited

Advances in Fire Suppression and Battery Safety: Li-Ion and VRFB | Sumitomo Electric. Accessed 17 Aug. 2026. https://sumitomoelectric.com/products/flow-batteries/stories/advances-in-fire-suppression-and-battery-safety-in-lithium-ion-and-flow-batteries.

Agusdinata, Datu Buyung, et al. "Socio-Environmental Impacts of Lithium Mineral Extraction: Towards a Research Agenda." Environmental Research Letters, vol. 13, no. 12, Nov. 2018, p. 123001. https://doi.org/10.1088/1748-9326/aae9b1.

Alotto, Piergiorgio, et al. "Redox Flow Batteries for the Storage of Renewable Energy: A Review." Renewable and Sustainable Energy Reviews, vol. 29, Jan. 2014, pp. 325–35. https://doi.org/10.1016/j.rser.2013.08.001.

Aurbach, D., et al. "Prototype Systems for Rechargeable Magnesium Batteries." Nature, vol. 407, no. 6805, Oct. 2000, pp. 724–27. https://doi.org/10.1038/35037553.

"Batteries - Why Lithium-Ion?" Apple, Accessed 17 Aug. 2026. https://www.apple.com/batteries/why-lithium-ion/.

Bella, Federico, et al. "An Overview on Anodes for Magnesium Batteries: Challenges towards a Promising Storage Solution for Renewables." Nanomaterials, vol. 11, no. 3, Mar. 2021, p. 810. https://doi.org/10.3390/nano11030810.

Darling, Robert M., et al. "Pathways to Low-Cost Electrochemical Energy Storage: A Comparison of Aqueous and Nonaqueous Flow Batteries." Energy & Environmental Science, vol. 7, no. 11, Nov. 2014, pp. 3459–77. https://doi.org/10.1039/c4ee02158d.

Gao, Tao, et al. "Enhancing the Reversibility of Mg/S Battery Chemistry through Li+ Mediation." Journal of the American Chemical Society, vol. 137, no. 38, Sept. 2015, pp. 12388–93. https://doi.org/10.1021/jacs.5b07820.

Glushenkov, Alexey M. "Recent Commentaries on the Expected Performance, Advantages and Applications of Sodium-Ion Batteries." Energy Materials, vol. 3, no. 2, Mar. 2023, p. N/A-N/A. https://doi.org/10.20517/energymater.2022.70.

Greenwood, Matthew, et al. "The Battery Component Readiness Level (BC-RL) Framework: A Technology-Specific Development Framework." Journal of Power Sources Advances, vol. 14, Mar. 2022, p. 100089. https://doi.org/10.1016/j.powera.2022.100089.

Greitemeier, Tim, et al. "China's Hold on the Lithium-Ion Battery Supply Chain: Prospects for Competitive Growth and Sovereign Control." Journal of Power Sources Advances, vol. 32, Apr. 2025, p. 100173. https://doi.org/10.1016/j.powera.2025.100173.

Johnson, Peter. "CATL Is Launching Sodium-Ion Batteries in EVs in 2026, Aiming for 370+ Miles Range." 22 Apr. 2026, https://electrek.co/2026/04/22/catl-launching-sodium-ion-batteries-evs-2026/.

Kelly, Jarod C., et al. "Energy, Greenhouse Gas, and Water Life Cycle Analysis of Lithium Carbonate and Lithium Hydroxide Monohydrate from Brine and Ore Resources and Their Use in Lithium Ion Battery Cathodes and Lithium Ion Batteries." Resources,

Conservation and Recycling, vol. 174, Nov. 2021, p. 105762. https://doi.org/10.1016/j.resconrec.2021.105762.

Leong, Kee Wah, et al. “Next-Generation Magnesium-Ion Batteries: The Quasi-Solid-State Approach to Multivalent Metal Ion Storage.” Science Advances, vol. 9, no. 32, Aug. 2023, p. eadh1181. https://doi.org/10.1126/sciadv.adh1181.

Leung, Adrian. “CATL Says Next-Gen Sodium-Ion Battery Supports 500 Km Range, Readies for 2026 Mass Production.” CarNewsChina.Com, 18 Sept. 2025, Accessed 17 Aug. 2026. https://carnewschina.com/2025/09/18/catl-says-next-gen-sodium-ion-battery-supports-500-km-range-readies-for-2026-mass-production/.

Li, Liyu, et al. “A Stable Vanadium Redox-Flow Battery with High Energy Density for Large-Scale Energy Storage.” Advanced Energy Materials, vol. 1, no. 3, 2011, pp. 394–400. https://doi.org/10.1002/aenm.201100008.

Li, Runze, et al. “Research on the Current Development Status of Redox Flow Batteries.” Molecules, vol. 31, no. 6, Jan. 2026, p. 943. https://doi.org/10.3390/molecules31060943.

Lyon, Peter. “These Are America’s 3 Best-Selling EVs Right Now, According To Car And Driver.” Forbes, 10 May 2026, Accessed 17 Aug. 2026. https://www.forbes.com/sites/peterlyon/2026/05/10/these-are-americas-three-best-selling-evs-right-now/.

Machín, Abniel, and Francisco Márquez. “Sodium-Ion Batteries: Advances, Challenges, and Roadmap to Commercialization.” Batteries, vol. 12, no. 4, Apr. 2026, p. 131. https://doi.org/10.3390/batteries12040131.

Malachi, Idowu O., et al. “Looking Beyond Lithium for Breakthroughs in Magnesium-Ion Batteries as Sustainable Solutions.” Sustainability, vol. 17, no. 9, Jan. 2025, p. 3782. https://doi.org/10.3390/su17093782.

Manthiram, Arumugam. “A Reflection on Lithium-Ion Battery Cathode Chemistry.” Nature Communications, vol. 11, no. 1, Mar. 2020, p. 1550. https://doi.org/10.1038/s41467-020-15355-0.

Matsui, Masaki, et al. “Destabilized Passivation Layer on Magnesium-Based Intermetallics as Potential Anode Active Materials for Magnesium Ion Batteries.” Frontiers in Chemistry, vol. 7, Jan. 2019. https://doi.org/10.3389/fchem.2019.00007.

Medina, Alejandro, et al. “Advancing towards a Practical Magnesium Ion Battery.” Materials, vol. 14, no. 23, Jan. 2021, p. 7488. https://doi.org/10.3390/ma14237488.

Mohtadi, Rana, and Fuminori Mizuno. “Magnesium Batteries: Current State of the Art, Issues and Future Perspectives.” Beilstein Journal of Nanotechnology, vol. 5, 2014, pp. 1291–311, https://doi.org/10.3762/bjnano.5.143.

Mousavian, Seyed Mohammad Hossein, and Vasily Anatolievich Bautin. “Rechargeable Magnesium Batteries: Overcoming Challenges for High-Efficiency Energy Applications.” Journal of Energy Storage, vol. 126, Aug. 2025, p. 117079. https://doi.org/10.1016/j.est.2025.117079.

"New Report Showcases How Innovation Can Fast Track Affordable Energy Storage." Energy.Gov, 1 Aug. 2024, Accessed 17 Aug. 2026. https://www.energy.gov/oe/articles/new-report-showcases-how-innovation-can-fast-track-affordable-energy-storage.

Nichols, Conrad. Redox Flow Batteries Market 2026-2036: Forecasts, Markets, Technologies and Players. 2025. Accessed 17 Aug. 2026. https://www.idtechex.com/en/research-report/redox-flow-batteries-market/1134.

Nolte, Oliver, et al. "Trust Is Good, Control Is Better: A Review on Monitoring and Characterization Techniques for Flow Battery Electrolytes." Materials Horizons, vol. 8, no. 7, July 2021, pp. 1866–925. https://doi.org/10.1039/d0mh01632b.

Olabi, Abdul Ghani, et al. "Redox Flow Batteries: Recent Development in Main Components, Emerging Technologies, Diagnostic Techniques, Large-Scale Applications, and Challenges and Barriers." Batteries, vol. 9, no. 8, Aug. 2023, p. 409. https://doi.org/10.3390/batteries9080409.

Orikasa, Yuki, et al. "High Energy Density Rechargeable Magnesium Battery Using Earth-Abundant and Non-Toxic Elements." Scientific Reports, vol. 4, no. 1, July 2014, p. 5622. https://doi.org/10.1038/srep05622.

Paul, Sandeep, et al. "Benchmarking the Performance of Lithium and Sodium-Ion Batteries With Different Electrode and Electrolyte Materials." Energy Storage, vol. 6, no. 7, 2024, p. e70068. https://doi.org/10.1002/est2.70068.

Phillips, Lawrence D. "Multi-Criteria Analysis: A Manual." ResearchGate, Accessed 17 Aug. 2026. https://www.researchgate.net/publication/30529921_Multi-Criteria_Analysis_A_Manual.

Pinto-Bautista, Sebastián, et al. "Prospective Life Cycle Assessment of an Electric Vehicle Equipped with a Model Magnesium Battery." Energy, Sustainability and Society, vol. 14, no. 1, July 2024, p. 44. https://doi.org/10.1186/s13705-024-00475-y.

Ponrouch, A., et al. "Multivalent Rechargeable Batteries." Energy Storage Materials, vol. 20, July 2019, pp. 253–62. https://doi.org/10.1016/j.ensm.2019.04.012.

Rasheev, Hristo, et al. "Redox Hyperactive MOF for Li+, Na+ and Mg2+ Storage." Molecules, vol. 27, no. 3, Jan. 2022, p. 586. https://doi.org/10.3390/molecules27030586.

Regulation - 2023/1542 - EN - Batteries Regulation - EUR-Lex. Accessed 17 Aug. 2026. https://eur-lex.europa.eu/eli/reg/2023/1542/oj/eng. Accessed 17 Aug. 2026.

Schirber, Michael. "Sodium as a Green Substitute for Lithium in Batteries." Physics, vol. 17, Apr. 2024, p. 73. https://doi.org/10.1103/Physics.19.s100.

Schirn, Alexandra. "IEEE 1725-2021: Rechargeable Batteries For Mobile Phones - ANSI Blog." The ANSI Blog, 2 Apr. 2024, Accessed 17 Aug. 2026. https://blog.ansi.org/ansi/ieee-1725-2021-rechargeable-batteries-mobile-phones/.

Skyllas-Kazacos, M., et al. "Progress in Flow Battery Research and Development." Journal of The Electrochemical Society, vol. 158, no. 8, June 2011, p. R55. https://doi.org/10.1149/1.3599565.

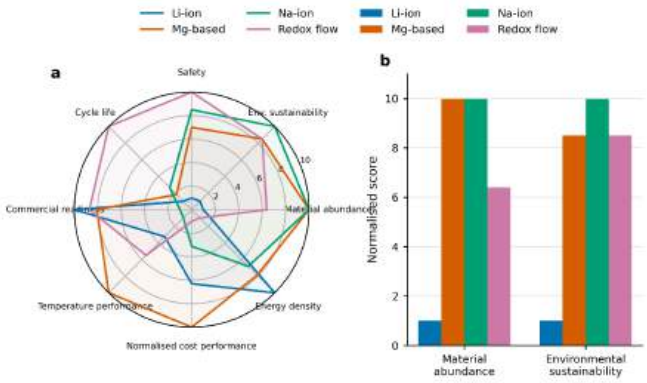
Song, Jaehee, et al. “Mapping the Challenges of Magnesium Battery.” The Journal of Physical Chemistry Letters, vol. 7, no. 9, Apr. 2016, pp. 1736–49. https://doi.org/10.1021/acs.jpclett.6b00384.
Tolmachev, Yuriy. “Flow Batteries From 1879 To 2022 And Beyond.” Qeios, Dec. 2022. https://doi.org/10.32388/G6G4EA.
Tricco, Andrea C., et al. “PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation.” Annals of Internal Medicine, vol. 169, no. 7, Oct. 2018, pp. 467–73, https://doi.org/10.7326/M18-0850.
Usiskin, Robert, et al. “Fundamentals, Status and Promise of Sodium-Based Batteries.” Nature Reviews Materials, vol. 6, no. 11, Nov. 2021, pp. 1020–35. https://doi.org/10.1038/s41578-021-00324-w.
Vera, María L., et al. “Environmental Impact of Direct Lithium Extraction from Brines.” Nature Reviews Earth & Environment, vol. 4, no. 3, Mar. 2023, pp. 149–65. https://doi.org/10.1038/s43017-022-00387-5.
Whittingham, M. S. “Electrical Energy Storage and Intercalation Chemistry.” Science, vol. 192, no. 4244, June 1976, pp. 1126–27. https://doi.org/10.1126/science.192.4244.1126.
Why Are Lithium-Ion Batteries, and Not Some Other Kind of Battery, Used in Electric Cars and Grid-Scale Energy Storage? | MIT Climate Portal. Accessed 17 Aug. 2026. https://climate.mit.edu/ask-mit/why-are-lithium-ion-batteries-and-not-some-other-kind-battery-used-electric-cars-and-grid. Accessed 17 Aug. 2026.
Winsberg, Jan, et al. “Redox-Flow Batteries: From Metals to Organic Redox-Active Materials.” Angewandte Chemie International Edition, vol. 56, no. 3, 2017, pp. 686–711. https://doi.org/10.1002/anie.201604925.
Wu, Fangshuai, et al. “Lithium Resources and Novel Strategies for Their Extraction and Purification.” Npj Materials Sustainability, vol. 3, no. 1, Sept. 2025, p. 30. https://doi.org/10.1038/s44296-025-00069-5.

Data availability

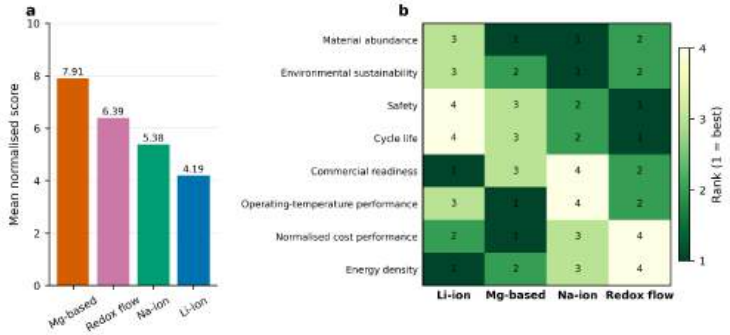
The normalized score matrix and analysis tables are summarised in the data table

Feature	Li	Mg	Na	Redox flow	Li-ion norm	Mg-ion norm	Na-ion norm	Redox norm
Energy density	200	160	140	30	10.0	7.9	6.8	1.0
Cycle life	4000	5000	6000	15000	1.0	1.8	2.6	10.0
Cost/kWh	110	75	140	160	6.3	10.0	3.1	1.0
Safety	7	9	9.5	10	1.0	7.0	8.5	10.0
Material abundance	5	10	10	8	1.0	10.0	10.0	6.4
Sustainability Score	6	8.5	9	8.5	1.0	8.5	10.0	8.5
Commercial readiness	10	8.5	3	9	10.0	8.1	1.0	8.7
Operating temp	80	110	70	90	3.3	10.0	1.0	5.5

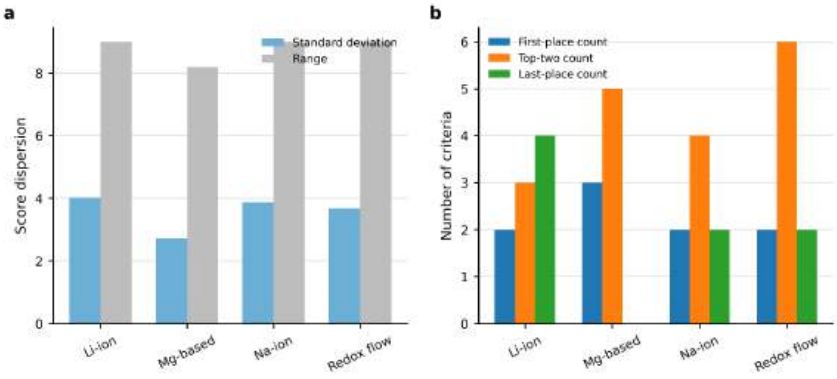
Supplementary Data



Supplementary Data Figure 1: Additional views of the criterion profiles. a, Radar representation of the eight normalised scores. b, Direct comparison of material abundance and environmental sustainability. Higher values indicate more favourable performance within the analysed dataset.



Supplementary Data Figure 2: Descriptive score summaries. a, Arithmetic mean of the eight normalised scores. b, Criterion-level rank, where 1 indicates the highest score and tied technologies retain the same rank.



Supplementary Data Figure 3: Score balance and criterion dominance. a, Standard deviation and range of the normalised scores; lower values indicate a more even profile but not necessarily a better technology. b, Counts of first-place, top-two and last-place criterion rankings. Ties are counted for each tied technology.

Agrivoltaics in California: Balancing Water, Energy, and Farmland in the Central Valley
By Angela Jin

Abstract

California's Central Valley is facing three interconnected crises. Groundwater depletion has increased rapidly since 2003 and some basins have overdraft rates nearly five times the 60-year average. At the same time, the state has started implementing large-scale solar with almost 50 percent of California's solar generated in the San Joaquin Valley, mostly on farmlands. Agriculture and solar energy development contend for the available land, raising a conflict of land use that can jeopardize agriculture, despite the growing use of green technologies. Agrivoltaics, which means mounting solar panels above agricultural crops, has become a possible solution to the conflict. This review summarizes the research on agrivoltaic systems in California and looks at the water-energy-land trade-offs associated with them. Agrivoltaics can lower irrigation needs by 20-40%. At the same time, it produces electricity at utility scale on a similar footprint. But crop reactions differ greatly because while tomatoes and some peppers yield more, crops like wheat, cucumbers, and potatoes can suffer yield losses due to excessive shading. Successful application of agrivoltaic technology lies primarily in selecting the right crops corresponding to solar panel density, topography and soil type. Geographic analyses show that installation of agrivoltaic technology is feasible in southern San Joaquin Valley owing to its high solar potential, extreme groundwater stress, existing conflicts and developed solar infrastructure. However, regions have obstacles from interconnection backlog, lack of financing, uncertain farmland preservation policy, and absence of any incentive programs for agrivoltaic implementation. Studies indicate that agrivoltaics can cut down water usage levels, as well as help conserve land for agriculture. The technology can be used to help solve both issues but it should be noted that its effectiveness is dependent on factors such as the technology's deployment in relation to certain geographic, agronomic, and policy conditions.

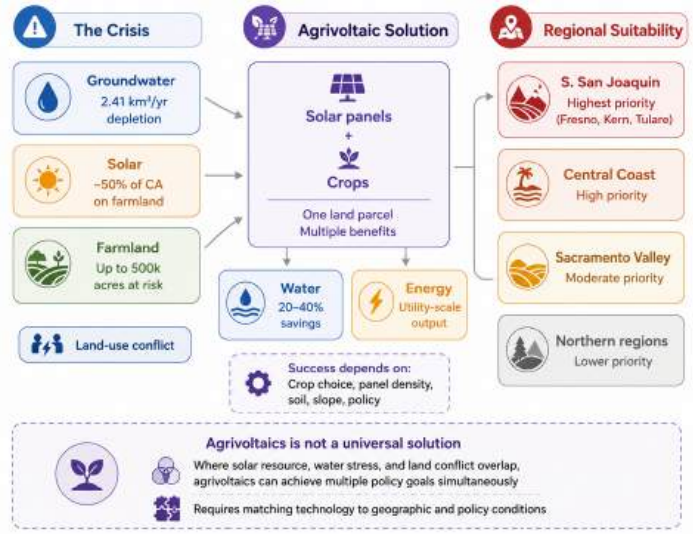


Fig 1: Graphical Abstract

1. Introduction

California is one of the most productive agricultural regions in the US and produces more than 400 commodities annually with a total output value of over $50 billion. Irrigation is mainly responsible for its production, as groundwater is a key source of water, and most of this production occurs in the Central Valley, where the San Joaquin Valley and Sacramento Valley are located. Groundwater provides 40 to 60 percent of California's irrigation supply in dry years and can account for as much as two-thirds of the Central Valley's irrigation supply during drought, making groundwater essential to Central Valley agriculture but also leading to overextraction ("Groundwater in California"; Liu et al.).

Repeated droughts have exposed how vulnerable this system is. Liu et al. analyzed two decades of NASA GRACE satellite observations and concluded that groundwater depletion in the Central Valley has been speeding up since 2003, from 1.86 cubic kilometers per year from 1961 to 2021 to 2.41 cubic kilometers per year from 2003 to 2021, and reaching a peak of 8.58 cubic kilometers per year during the megadrought of 2019-2021, almost five times above the 60-year average (Liu et al.). The recovery has not been full: about a third of what was lost in the 2007–2009 drought and about a fifth of what was lost in the 2012–2016 drought have been recovered (Liu et al.). California enacted the Sustainable Groundwater Management Act (SGMA) in 2014 to address long-term groundwater overdraft, but as of 2026 the state still identifies 21 groundwater basins as critically overdrafted, resulting in damage that could be permanent, such as aquifer compaction in the Sacramento Valley ("Groundwater in California"). SGMA gives most authority to over 260 Groundwater Sustainability Agencies in local areas and imposes stricter deadlines on badly overstressed basins, requiring them to become sustainable by 2040 instead of 2042, which matters here since most San Joaquin Valley basins discussed in this paper fall into that tougher category ("SGMA Explainer").

At the same time, the state has ambitious renewable energy goals and utility-scale solar is moving fast, putting a second layer of stress on the same land. Both solar farms and agriculture need open, sunny land close to existing infrastructure, which is common across the Central Valley's farmland. Ag land is already short on water resources, and traditional solar energy competes with agriculture for space, producing a land-use dilemma in which renewable energy may reduce the farmland available for agriculture.

The dispute is not hypothetical. Almost 50% of California's 100,000 acres of solar farms are located in the San Joaquin Valley and most have been built over agricultural land affected by SGMA (Stid et al., "Solar Array Placement"). An inventory of the Central Valley using remote sensing has found that there are 1,006 ground-mounted solar arrays installed from 2008 to 2018, of which 694 or around 69 percent have been co-located with agriculture and which provide a combined capacity of 2.1 gigawatts, showing that many of them used poor spacing and placement, providing evidence that the land-use conflict that this review will discuss is not hypothetical but rather visible in what solar looks like today (Stid et al., "Solar Array Placement"). As of 2023, Kern County has either approved the construction of or was already constructing over 3,000 megawatts of solar power capacity, some of which was either on

previously-utilized land or active farmland ("Solar Energy and Groundwater"). Westlands Water District has suggested turning 136,000 acres of drought-affected farmland into a solar energy and battery storage initiative, framing it as a means of countering declining groundwater rather than energy demand alone ("More Farmland Going Solar"). Meanwhile, separate modeling suggests that meeting SGMA's groundwater balance requirements could require taking up to 500,000 acres of irrigated San Joaquin Valley farmland out of production, close to 10 percent of the state's irrigated land, regardless of how solar development proceeds ("Solar Energy and Groundwater"). Agrivoltaics tackles the question of whether a single piece of land could take care of water conservation, solar development and farmland conservation all at once.

Agrivoltaics, the combination of solar panels and crop cultivation in a single location, has been suggested as a solution to this dilemma. This method can provide advantages for agricultural production, including lower drought stress, better water use, more food grown, preservation of agricultural lands, and solar energy generation at the same time ("AGRIVOLTAICS"). Initial field studies locating solar facilities on farms have observed changes in microclimate, panel temperature, and soil moisture, and identified interactions between crop production and solar infrastructure (Barron-Gafford et al.). According to a 2025 Nature Reviews Clean Technology review, global deployment of agrivoltaic systems could produce 66 to 385 petawatt-hours of electricity annually from appropriate sites, though installation costs remain 20 to 90 percent higher than conventional systems ("Scientific Frontiers of Agrivoltaic Cropping Systems"). Research on agrivoltaics' compatibility with California is still inadequate, since most existing literature highlights only a single benefit, like yield, energy output, or water savings, rather than linking energy, water, and land use as an interconnected system.

This review integrates existing research to analyze the water-energy-land trade-offs of agrivoltaics, evaluate the conditions its success or failure depends on, and identify which parts of California and which policy environments are most likely to make it a genuinely sustainable solution. The goal is to determine whether agrivoltaics can help address California's drought, energy, and farmland pressures together, rather than treating it as a promising technology in the abstract.

2. How Agrivoltaics Works, and the Productivity-Water-Energy Triangle

2.1 The Mechanism: Shading, Microclimate, and Soil

Agrivoltaic systems install solar panels at a height of several meters on posts so that crops can grow underneath while the panels generate electricity. Shade from the panels can help avoid heat stress and maintain soil humidity, but too much shade inhibits crop growth ("Influence of Agrivoltaic System-Induced Microclimate Modifications"). Field research conducted beneath solar panels has found wind speed reductions of 27.6 to 42.3 percent, along with a temperature increase of about 2.3 degrees Celsius and a humidity increase of nearly 36 percent, and a long-term study in orchards found soil humidity up around 11.8 percent and soil temperature

down 0.8 degrees Celsius compared to unprotected soil ("An Observational Study"; "Data-Driven Analysis of the Effect of Agrivoltaics Systems on Soil and Air Conditions"). A three-dimensional computational fluid dynamics model of a vertical agrivoltaic system, validated against experimental data, found that photovoltaic panel shading resulted in a decrease of solar radiation reaching the land area by 38% while likewise reducing wind speed on the ground level, confirming the field trial findings above ("Computational Fluid Dynamics Modelling").

Soil moisture retention is the most relevant topic for California's water crisis. A controlled experiment on a pasture at Oregon State University found soil moisture levels much greater under solar panels during the growing season, 90 percent more biomass produced in the late season, and water efficiency increased by 328 percent compared with unshaded areas (Hassanpour Adeh et al.). Studies in France showed that panel shading reduces irrigation by 14 to 29 percent compared to growing without shade (Marrou et al.). A 2024 systematic review of 33 field studies confirmed this as a broad pattern, finding that agrivoltaic systems improved water-use efficiency by 20 to 47 percent and lowered air and soil temperatures by 1 to 4°C across a range of climates ("Impacts of Agrivoltaic Systems on Microclimate"). Recent research has begun testing whether water savings can be engineered rather than left as a side effect of shading: a 2026 Spanish study combined shaded agrivoltaic technology with deficit irrigation, which restricts water use during crop growth, and reduced irrigation by 50 percent while preserving land-use efficiency ("New Irrigation Tech").

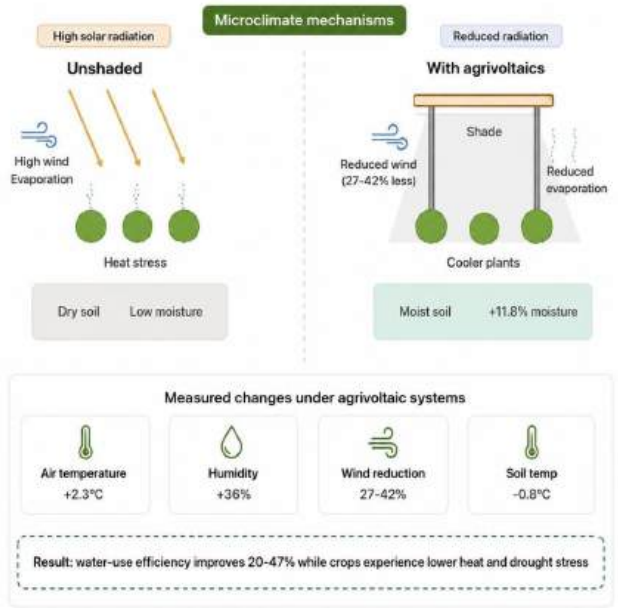


Fig 2: Microclimate mechanisms diagram showing measured changes under agrivoltaic system

2.2 The Triangle: Productivity, Water Conservation, and Energy Generation

This creates a design dilemma in agrivoltaics modeled as a triangle having three competing objectives at its ends: crop productivity, water conservation, and energy generation. More coverage by panels means higher electricity from an acre but also lower light for the crops,

and beyond a certain point, yield losses will be higher than its water and heat-stress solutions. Research has shown that shading from panels and other changes to microenvironment conditions can often result in decreased photosynthetically active radiation, negatively affecting the efficiency of photosynthesis and producing problems with crop yield and quality if shading is too great ("Agrivoltaics Revisited"). A German field trial that grew celeriac, winter wheat, potato, and grass-clover in rotation near Lake Constance found that the panels reduced photosynthetically active radiation reaching the crops by about 30 percent on average, making PAR reduction the most consistent constraint on yield across all four crops, even though the panels also measurably reduced soil and air temperature and, for some crops in some seasons, increased soil moisture ("Agrivoltaic System Impacts on Microclimate"). The mechanism that protects water resources can also suppress crop yield, so it is impossible to maximize all three goals at once, which means these goals need to be balanced through choices like spacing, panel number, height, and position.

Row spacing illustrates how sensitive this balance can be. An optimization model of vertically mounted, bifacial agrivoltaic systems found that reducing row distance from 20 meters to 5 meters, a change made to increase panel density and therefore electricity output per acre, reduced the crop yield of potato and oat by about 50 percent, even as it improved the land equivalent ratio above 1.2, meaning the combined system still out-produced separate mono-use land overall despite the steep per-crop yield cost ("Optimization of Vertically Mounted Agrivoltaic Systems"). The productivity-water-energy triangle serves as an engineering idea instead of just an abstract concept, since a decision that might optimize energy production in theory might reduce agricultural yield by 50% in practice.

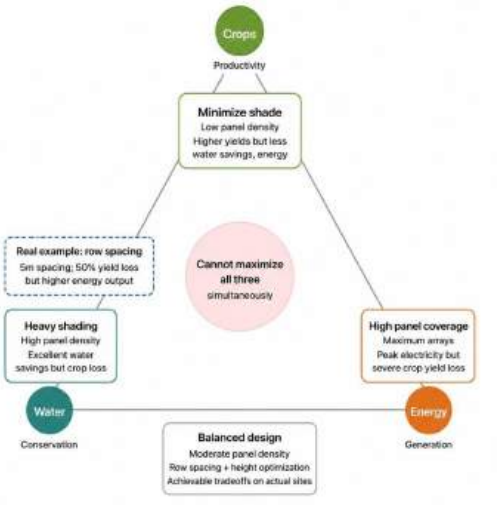


Fig 3: The productivity-water-energy triangle diagram

2.3 Crop-Specific Responses

Research shows that plant response to light interference varies by crop. A survey by the National Renewable Energy Laboratory of agrivoltaic research sites across the country found

that average tomato output was twice as much as at non-agrivoltaic sites, while wheat, cucumber, potato, and lettuce were negatively affected ("Cool Tomatoes?"). This crop-focused type of agrivoltaics remains a minority option nationally: NREL's own research network statistics show pollinator and native habitat installations comprise the largest number, at over 400, versus only 35 installations for crop production, so California is at the cutting edge of applying this technology to field crops ("Lighting the Way for Agrivoltaics"). The tomato outcome seems linked to heat stress: when temperatures get too high, tomato plants stop producing fruit from flowers, and the cooling effect of solar panels may counteract this, an important observation for California, where the bulk of American tomatoes are grown ("Cool Tomatoes?"). In a study of chiltepin pepper, jalapeno, and cherry tomato in a desert climate, total chiltepin fruit production was three times greater with agrivoltaics, jalapeno water-use efficiency increased by 157 percent, and cherry tomato water-use efficiency increased by 65 percent, with fruit production doubled (Barron-Gafford et al.; "Benefits of Agrivoltaics").

Leafy greens act differently in terms of growth, rather than just gaining or losing yield. Research shows that lettuce grown under agrivoltaic systems can compensate for a 50 percent shortage of light by making larger amounts of total leaf area per plant, thus keeping its yield and interception of radiation (Marrou et al.). A Turkish field trial, growing lettuce under a photovoltaic structure built from repurposed pruning residue rather than conventional steel, found that full irrigation produced the best marketable head weight and color of any tested watering level under agrivoltaic shade ("Agrivoltaic Impact on Some Lettuce Quality Attributes"). A systematic review of 54 studies showed a high degree of variability across studies regarding yields for both lettuce and tomato, the two crops most studied in agrivoltaics, indicating that local design factors are crucial to the outcome and not the crop itself (Widmer et al.). Zucchini trials demonstrate that even within a single crop, a trade-off between productivity and water exists: high shade reduces temperature and drought stress and improves plant functioning in general, while at the same time detracting from reproductive yield ("High-Shade Dryland Agrivoltaic Conditions"). A trial growing potatoes beneath, between, and outside simulated solar panels found a similar non-linear pattern: potatoes in the partially shaded zone between panels produced the highest total yield of the three conditions tested, suggesting that for at least some root crops, agrivoltaic row spacing creating intermittent rather than constant shade may outperform full sun ("The Effect of Agrivoltaic System on Nutrient Content"). A Purdue University study with maize, a C4 crop not renowned for shade tolerance, found that with roughly 20 to 25 percent of shade, the average yield loss for a maize hybrid is only 7.7 percent, partly due to the safety net from soil moisture ("Multi-Year Study of Maize"). In a two-year field experiment in Denmark with wheat, grass-clover, and soybean under different agrivoltaic panel setups, wheat output remained stable during both years and protein and gluten content even rose compared to the open control, though soybean yield fell by about one-third, which the researchers trace to differences in canopy and root architecture between the two crops under reduced wind speed ("Agrivoltaic Systems: Trade-Offs").

Overall, these findings lead to a working principle for the rest of this review: agrivoltaics

succeed not as a property of the technology alone, but by fitting a crop's shade tolerance to a specific panel density and climate.

3. Effects on Water Use

3.1 Evapotranspiration and Irrigation Demand

Evapotranspiration, the loss of water from the soil and plant, is one of the important factors for the reduced amount of water needed by agrivoltaic farms. The Penman-Monteith equation was used to show that under arid and semi-arid conditions, agrivoltaics helps save water depending on coverage, between 30 and 40 percent ("Agrivoltaic System Design Tools"). A comprehensive review of agrivoltaic studies found increases in water-use efficiency metrics of between 150 and 300 percent, along with a decrease in average irrigation requirement of about 14 percent, with the biggest improvements in arid and semi-arid regions ("A Systematic Review of Agrivoltaics: Productivity, Profitability"). A study from France found cumulative evapotranspiration in an agrivoltaic lettuce system to be 142 millimeters, or 59 percent, less than in conventional operations, reflecting 289 liters in irrigation savings ("A Systematic Review of Agrivoltaics: Productivity, Profitability").

These savings are not limited to green vegetables. Mean evapotranspiration savings from a maize research study reached 46 mm for agrivoltaic methods compared to traditional farming, and several soil evaporation tests reported savings of 14 percent to 33 percent due to evaporation being blocked by panels ("A Systematic Review of Agrivoltaics: Productivity, Profitability"; "Water Evaporation Reduction").

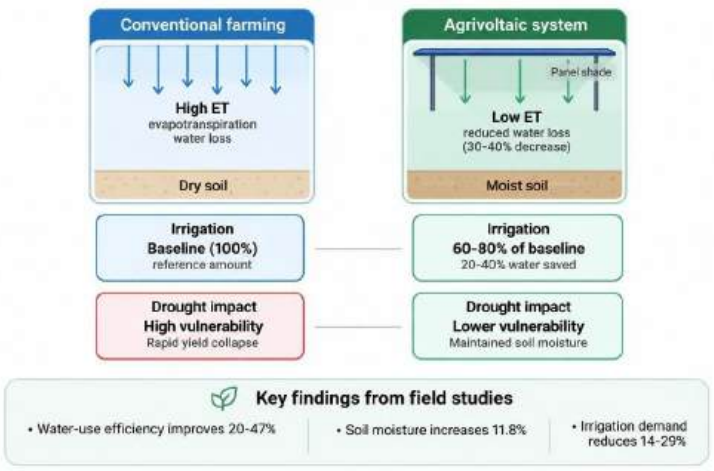


Fig 4: Water savings through agrivoltaics diagram comparing conventional farming to agrivoltaic systems

3.2 Drought Resilience

Agrivoltaics not only aids with regular irrigation but also shields crops from extreme drought stress, a significant advantage for California where years of drought threaten Central Valley agriculture more than average climatic conditions. Research saw that drought conditions

led agrivoltaic crops to outperform crops grown through traditional farming practices, while agrivoltaics slowed down soil water loss and allowed the moisture to stay longer–something that could lead to greater resistance to drought ("Impact of Agrivoltaics Irrigation and Soil Moisture Dynamics"). This pattern held even for crops not typically associated with shade tolerance: rainfed sugar beet yields under agrivoltaics improved significantly during hot, dry seasons with limited rainfall, closing much of the gap with conventional agriculture's yield under water-stressed conditions ("Impact of Agrivoltaics Irrigation and Soil Moisture Dynamics"). Field research from East Africa indicates there is yet another drought-resilience path, wherein rainwater directly harvested from photovoltaic panel runoff was applied to augment irrigation at agrivoltaic sites in Tanzania and Kenya, decreasing the need for groundwater or surface water during droughts ("Harvesting the Sun Twice"). Runoff harvesting is not seen in most California agrivoltaic proposals so far, but the East Africa findings show that it is useful for constructing drought-resilient installations.

This is relevant to the California droughts mentioned above. If agrivoltaic soil moisture retention works as the studies conclude, it is not merely a matter of less water used, but also of managing dry years better, which matters from a policy point of view because it suggests agrivoltaics might lessen damage to farmer yield and revenue during prolonged droughts rather than only providing modest savings in ordinary years.

3.3 Can Agrivoltaics Actually Solve Drought?

It is important to be clear about what the evidence does and does not say. Although agrivoltaics decrease water consumption per acre and increase drought resiliency, no studies indicate it puts water back into drained aquifers or creates new water sources. Freshwater withdrawals already exceed 100 percent of renewable water resources in parts of the southwestern United States, and the region's more than 8 million irrigated acres face a projected 10 to 20 percent decline in future water supply, so agrivoltaics functions as a demand-side tool: it reduces how much water each acre needs rather than solving the supply shortfall driving SGMA-mandated cuts ("The Long-Term Suitability of Agrivoltaics"). Agrivoltaics is more accurately a technique that helps a limited water supply last longer and mitigates drought years, rather than a solution to drought and groundwater depletion themselves.

4. Why California, and Where Agrivoltaics Actually Works

4.1 California's Structural Fit for Agrivoltaics

California has a unique combination of three factors. The first is solar resource: California has more installed solar power than any other state, combined with strong solar irradiance. The second is the drought-pressure situation already discussed above, which is reducing agricultural water availability regardless of agrivoltaics. The third is land-use conflict, where solar developers and agriculture companies compete for the same land, as seen in Kern County and the Westlands Water District (Stid et al., "Solar Array Placement"; "More Farmland

Going Solar"). No other U.S. agricultural region combines this level of solar potential, groundwater stress, and land-use conflict, which is why California is the right place to research and formulate agrivoltaics policy.

4.2 Cost and Installation Factors

Agrivoltaics has a higher installation cost than typical ground-mounted solar, and this price difference depends heavily on design specifics such as panel height. The steel, complexity, and installation difficulty of higher agrivoltaic systems make them more expensive than installations meant for crops grown with hand labor ("Agrivoltaic Solar Mounting"). Elevated mounting structures with necessary machinery clearance usually raise installation costs by about 20 to 40 percent compared to fixed-tilt ground-mounted systems ("Agrivoltaic Solar Mounting"). Broader international data show cost increases specifically related to agrivoltaics ranging from about 4 to 115 percent, depending on agricultural activity, panel height, and structural configuration (Trommsdorff et al.). This agrivoltaic premium is an important part of overall cost, causing difficulties for smaller companies without major investors.

Federal and state incentive programs currently offset much of this gap. The federal Investment Tax Credit covers 30 percent of installation costs, and USDA's Rural Energy for America Program offers grants covering up to an additional 25 percent for qualified agricultural producers, together capable of offsetting 50 to 75 percent of initial capital costs ("The Use and Potential of Agrivoltaics"). Federal solar incentives have been somewhat unpredictable recently, but evidence from the solar industry shows policy instability hinders investment while deployment increases significantly when incentives are stable and long-term, suggesting the cost of agrivoltaics in California will depend greatly on incentive stability ("United States Agrivoltaics Policy Landscape").

4.3 Land-Use Regulation and Farmer Acceptance

A lot of farmers doubt agrivoltaics because of tough permits and zoning processes, unsure of how profitable it will be, and lacking knowledge on growing new crops under panels ("Drivers of Agrivoltaic Perception"). California has its own farmland-protection system that complicates matters even more. Cities and counties have had the authority since 1965, under the Williamson Act, to restrict land for agricultural purposes in return for lower property taxes based on the valuation of the land. Those seeking to construct solar power facilities on Williamson Act land must either allow the contract to fall into non-renewal status, which takes nine years, or request that it be canceled through a petition that costs 12.5 percent of the unrestricted fair market value of the land ("The Williamson Act"; "Williamson Act Contract Removal"). Legal advice available for landowners participating in this process suggests that the option of nonrenewal, although slower, eliminates the cancellation fee altogether, leading a farmer considering the transition to agrivoltaics to confront a clear dilemma of speed versus cost ("Williamson Act in California"). California authorized a separate solar-use easement mechanism in 2011 that allows a Williamson Act contract to be rescinded and replaced with a solar-specific

easement, but this still requires state Department of Conservation approval and remains a slower, more uncertain path than converting unrestricted land ("The Williamson Act"). California currently lacks an agrivoltaics-specific incentive or streamlined permitting pathway comparable to those adopted elsewhere; Massachusetts, for instance, enacted a feed-in tariff adder specifically for agrivoltaic projects, and New Jersey authorized an agrivoltaics pilot program on farmland not already under a conservation easement ("United States Agrivoltaics Policy Landscape").

Economic modeling indicates that the crop chosen for growth has a great impact on the success of the agrivoltaic project. A 25-year Monte Carlo simulation of small farm agrivoltaic systems carried out on Kentucky farms growing tomatoes and strawberries demonstrates that profitability is driven by crop prices, so high-value crops like strawberries can allow for a significantly shorter payback period and increased overall returns compared to lower-value crops ("The Techno-Economic Viability of Agrivoltaics"). An undergraduate suitability study in Central California, using its own GIS model, reached a similar conclusion, marking tomatoes as the crop offering maximum potential for high revenue and good agrivoltaic efficiency, consistent with the earlier heat-stress findings (Camba).

4.4 Where Agrivoltaics Works Best

Suitability is a technical engineering consideration, not just a case of policy and climate. The most persistent limitation is slope, which plays a major role in examining sites. According to a review of 30 studies on agrivoltaic site feasibility by the Federal University of Rio de Janeiro, acceptable slope limits for ground-mounted systems may vary from 3 percent to 70 percent depending on design, but 63 percent of the studies set the slope threshold below 15 percent, favoring flatter areas with better agricultural potential and lower erosion risk ("The Impact of Land Slope"). A companion framework built around soil and terrain classification found that slopes below 15 percent represent the most balanced condition for dual agricultural-energy land use, slopes between 15 and 30 percent represent the maximum feasible range depending on crop type and erosion susceptibility, and above 20 percent, land may become more suitable for pasture or forestry than for the row crops agrivoltaic systems are generally designed around ("Soil Suitability and Technical-Construction Criteria"). This matters directly for California siting: the flat, low-elevation floor of the southern San Joaquin Valley may score well on these criteria, while the hillside vineyards of the Central Coast or the mountain-adjacent margins of the Sacramento Valley may face structural limits that no amount of water stress or solar resource can offset.

A second, more technical consideration relates to the soil's load-bearing properties and foundations. Conventional agrivoltaic systems can use driven steel poles instead of concrete footers since they set up more quickly, cause less soil disturbance, and can be removed when the project ends, leaving the land suitable for farming ("Agrivoltaic Solar Mounting"). Agricultural soils that have been heavily tilled or are subject to seasonal saturation, common in parts of the Central Valley, may need unique foundation designs, while sites with shallow bedrock or

low-bearing capacity soils may require the costlier concrete pier option ("Agrivoltaic Solar Mounting").

Panel height and row spacing are not fixed limits but design choices that depend on the equipment used on the farm underneath the array. Higher agrivoltaic structures require greater elevation and wider row spacing than usual ground-mounted solar installations, especially where a tractor or other farm equipment must operate under the panels ("Solar Farm Design Guide 2026"). Taken together, these findings mean two locations with the same climate, water stress, and grid distance can still differ substantially in viability once slope, soil bearing capacity, and equipment-access requirements are considered.

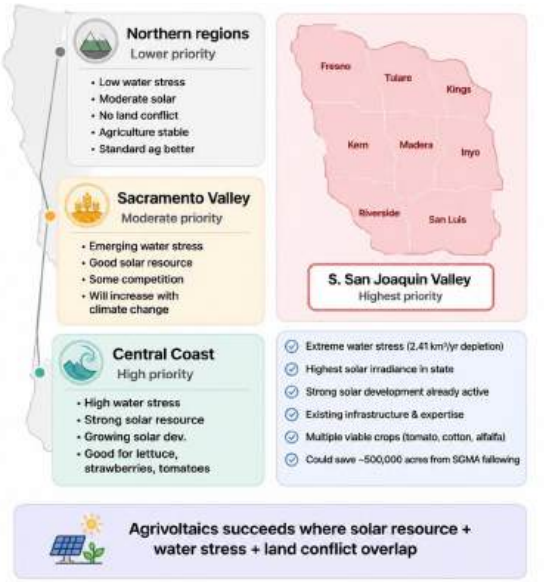


Fig 5: Regional suitability across California map showing priority tiers

5. Trade-Offs, Limitations, and Conditions for Success

5.1 Synthesizing the Trade-Offs Across This Review

The sections above have covered agrivoltaics under different perspectives: method, water savings, and land selection. Analyzed together, these three perspectives form a complicated picture where any perk is accompanied by a disadvantage, consistent with the productivity-water-energy triangle: every design solution comes with its own consequences. Installing more panels increases output and water savings but limits growth of shade-sensitive plants. Elevated mounting allows farmers to access agricultural machinery but incurs additional costs of about 20-40 percent over conventional installations ("Agrivoltaic Solar Mounting"). Building on flat ground makes construction easier but could require longer interconnection runs. Agrivoltaics does not resolve the dilemma of land usage; it transfers that dilemma into a series of engineering and policy decisions about a single piece of land.

5.2 Economic Viability Beyond Installation Cost

Installation cost is just the first hurdle in terms of finance a project has to overcome. Financing structure is another issue: many California farmers, especially those farming mid-sized or family operations rather than large corporate farms, do not have the financial means to withstand getting payments after many years; getting access to financing and not just the technology costs is more likely to draw the line between the farms that decide to implement agrivoltaics and those that do not ("The Techno-Economic Viability of Agrivoltaics"). The Central Valley agrisolar analysis presented earlier provides an interesting contrast: among the 925 systems analyzed through that study, the economic benefits enjoyed by farmers after switching from agriculture to solar power usually exceeded the drawbacks of lost agricultural output, suggesting that when there are no financial or grid access limitations, the economics of switching may still favor farmers even without factoring in agrivoltaic crop revenues (Stid et al., "Impacts of Agrisolar Co-Location"). Farmers entering into agrivoltaics are effectively placing bets on two different markets, the crop market and the electricity or power purchase agreement market; crop price is a key determinant of profitability in small-farm modeling, but electricity price has its own risk, as demonstrated by spring dips in California's wholesale solar prices due to oversupply in sunny weather ("The Techno-Economic Viability of Agrivoltaics"). A hydro-economic modeling study of the Central Valley raises a further complication: solar and agrivoltaic development can meaningfully aid local groundwater conservation, but only if paired with genuine changes in regional crop patterns, since concentrated conversion of land away from traditional crop production risks a form of "water leakage," in which farmers elsewhere in the valley respond to reduced regional supply by intensifying irrigation on their own remaining acreage, potentially offsetting some of the water savings the conversion was meant to produce (Yao et al.).

5.3 Grid Interconnection as an Underexamined Constraint

A majority of agrivoltaics models usually considered proximity to electric substations as a factor among other variables, but proximity may be necessary without being sufficient on its own. The average project completed in 2022 spent five years waiting in the interconnection queues, compared to three years in 2015 and almost two years in 2008, with total capacity waiting in the queues growing much faster than the amount of capacity actually completed ("US Grid Interconnection Backlog"; Lawrence Berkeley National Laboratory). This matters for the siting priorities presented here: the southern San Joaquin Valley counties rated highest priority have developed a lot of solar energy already, meaning the local grids may be closer to saturation than counties that have not developed as much. A parcel can pass all the suitability criteria but still experience a multi-year interconnection delay, reshaping its priority ranking.

5.4 Environmental Trade-Offs Beyond Water and Yield

One aspect that needs to be looked into is habitat and soil disturbance during construction. Even with pile-driven systems, which disturb less soil than poured concrete piers,

there is still a need for access roads, trenching, and short-term ground compaction ("Agrivoltaic Solar Mounting"). Pollination and biodiversity impacts should be considered as well, where agrivoltaics are expected to do better than conventional monoculture or free-standing solar systems. Research by NREL's InSPIRE program shows native prairie plants can be grown under solar panels, benefiting soil and providing pollinator habitat, though it usually takes three to four years to fully develop and some species take up to six years to appear ("Utility-Scale Solar Fields"). Separate trials on native seed mixes at utility-scale solar sites in Minnesota found that establishing pollinator habitat led to markedly more pollinator visits than unmanaged sites, though outcomes varied by seed mix and site conditions ("Testing Native Pollinator Habitat Seed Mix Establishment"). It is still unknown whether this applies to row-crop processes common in California's Central Valley, since most of this research has been conducted on pasture and native-vegetation systems. Central Valley soils that have experienced excessive plowing or temporary saturation require unique foundation treatments, but research on their longevity is limited by a lack of longitudinal data on how climate and machinery affect soil structure and fertility over time.

5.5 When Agrivoltaics Does Not Work

Steep or irregular terrain is the clearest disqualifying condition: most feasibility studies set slope thresholds below 15 percent, and land above roughly 20 percent slope shifts toward pasture or forestry suitability rather than the row-crop systems agrivoltaics is generally designed around ("The Impact of Land Slope"; "Soil Suitability and Technical-Construction Criteria"). Row crops that do not tolerate shade can also have adverse effects on yield; certain crops (e.g., wheat, cucumbers, potatoes, lettuce) yielded negatively under certain trial conditions, and if a farm relies heavily on one or more of these crops, the negative impact of considerable panel coverage may outweigh the gains made from decreased water expenses or rent income ("Cool Tomatoes?"). A third condition: where the land use conflict this review describes is absent from a particular parcel, agrivoltaics represents a real cost of solving a problem the land does not have, and regular agriculture or regular ground-mounted solar will likely be the better option.
Sites suffering from long interconnection delays or farms lacking access to funds or stable federal incentives might find agrivoltaics technically and agronomically feasible but impractical given the current market and policy landscape. Unlike the conditions in the previous paragraph, this is not a case of agrivoltaics being unsuited to the land's characteristics, but of financial and regulatory conditions not yet being in place to support the pace of implementation that California's land and water pressures would require.

5.6 Toward a Conditions-of-Success Framework

The agrivoltaic success case in California is not merely a technological issue but a combination of factors, including crops, weather, slope, soil, grid accessibility, and the financial aspects of each piece of land evaluated. A helpful way to evaluate sites and proposals is to ask four questions in order: whether the site meets the physical limits, whether the crop under

consideration can grow in the shade, whether grid interconnection is compatible with project financing, and whether site regulations are reasonable without factoring in federal help. This framework does not change the conclusion that the southern San Joaquin Valley is the best fit for the criteria, but it clarifies why that is a conclusion rather than an absolute fact.

6. Conclusion: Where Agrivoltaics Works, and What's Next for California

6.1 Agrivoltaics Is a Real Solution, But Not a Universal One

The findings throughout this paper do not endorse agrivoltaics as a universal solution to California's present water-energy-land challenge. The combination of agricultural activity and solar development lessens evapotranspiration, decreases irrigation requirements by 20-40 percent depending on crop and climate, and improves drought resistance. Nevertheless, it also increases installation costs by 20 to 115 percent depending on method, depends on matching crops to the right amount of shade, and does not add water to the aquifer. Productivity, water conservation, and energy generation form a genuine triangle: moving toward one goal detracts from another, and no single design maximizes all three at once. The conclusion, then, is not that agrivoltaics works, but that it works under certain conditions, namely geographic ones.

6.2 Ranking Suitability: Where the Data Points

Two independent GIS suitability studies reaching similar findings give more confidence in the ranking, since it does not depend on assumptions from a single model. The first study determined which regions of California farmland score best according to five criteria: crop type, farmland productivity, water stress, insolation, and proximity to electrical substations, concluding that southern Central Valley, Central Coast, and Southern California hold the most area of high agrivoltaic suitability, with Fresno, Tulare, Kern, Kings, Riverside, and Madera counties containing the most highly suitable land (Bihari). A second study on groundwater overdraft used a multi-criteria decision-making model to conclude that more than 2 million acres of active California farmland would be suitable for agrivoltaics across the state's 15 sub-basins, and that adopting agrivoltaics on that suitable land could save more than half of the Central Valley's groundwater overdraft ("Agrivoltaics for Managing Groundwater Overdraft"). A third regional study extending across the broader Southwest found that Arizona, New Mexico, and California contain the largest areas of most-suitable agrivoltaic land, with Imperial County consistently ranking among the most suitable counties, and that suitability in California and New Mexico is projected to increase further as climate change progresses ("The Long-Term Suitability of Agrivoltaics"). Based on this convergence, this review proposes a suitability hierarchy for California, synthesized from the three studies above rather than any single official source. The highest-priority deployment sites are in the southern San Joaquin Valley, including Fresno, Tulare, Kern, Kings, and Madera counties, which combine high solar output, high water stress, existing groundwater overuse, and existing deployment infrastructure and grower familiarity with solar leasing. A high priority tier includes Salinas Valley, the Central Coast, and

Imperial County, which have high suitability ratings but fewer land-conflict problems than the southern San Joaquin Valley, making them potential second-wave locations. A moderate priority tier includes sub-basins of the Sacramento Valley, which currently have lower water depletion than the San Joaquin Valley but are likely to approach high-priority status as climate change progresses. A lower priority tier includes agriculture in northern mountainous regions, where neither groundwater stress nor land-use conflict is significant and standard agriculture may remain the more efficient solution.

6.3 What This Ranking Means: Suitability Is Convergence, Not One Factor

The pattern that emerges from all three studies is not one criterion but the overlap of several. High insolation alone does not mean a site is suitable for agricultural use; the Mojave Desert has great solar potential but is not a significant agricultural area. High water stress does not imply suitability either, since some water-stressed land lacks the right crop-to-shade ratio or is not economically close to grid infrastructure. This is the practical version of the productivity-water-energy triangle described earlier: suitability is a multi-variable optimization problem, and the southern San Joaquin Valley scores well because it performs adequately across all five criteria simultaneously rather than excelling in any single one.

6.4 Outlook: What Would Make This Scale

The future of agrivoltaics in California is determined less by how effective the technology is, already verified in numerous studies, and more by government progress in utilizing the identified potential. This review underlines three specific gaps. First, incentive stability: federal incentive inconsistency is already a major obstacle to agrivoltaics investment nationally, and California lacks state-level agrivoltaics-specific incentives comparable to Massachusetts's feed-in tariff adder or New Jersey's pilot program ("United States Agrivoltaics Policy Landscape"). Second, policy integration: SGMA-driven fallowing and utility-scale solar siting are implemented as separate policy tracks, even though the suitability maps in this review show the two can occur on the same acres and sub-basins; a California program that explicitly directs fallowed, water-stressed SGMA land toward agrivoltaics could convert a potential 500,000-acre farmland loss into a dual-use asset instead of a pure loss. Third, crop-specific research investment: this review documented enormous variance in crop response, with tomato yields roughly doubling under agrivoltaics in national trials while wheat, cucumbers, potatoes, and lettuce showed negative impacts in the same survey. The specialty-crop economy of tomatoes, lettuce, and strawberries in California coincides well with the crops already researched extensively in agrivoltaic studies, giving the state a strong advantage.

Agrivoltaics does not work everywhere in California, nor does this review suggest that it does. But where solar resource, water stress, and land conflict already overlap, as in the southern San Joaquin Valley, the evidence points to a rare outcome in resource policy: achieving several conflicting goals at the same time.

Works Cited

"A Systematic Review of Agrivoltaics: Productivity, Profitability, and Environmental Co-Benefits." ScienceDirect, 8 Mar. 2025. Accessed May 12, 2026. www.sciencedirect.com/science/article/pii/S2352550925000569.

"Agrivoltaic Impact on Some Lettuce Quality Attributes and Photovoltaic Power Generation." Plants, National Center for Biotechnology Information, 2025. Accessed July 3, 2026. www.ncbi.nlm.nih.gov/pmc/articles/PMC12473297/.

"Agrivoltaic Solar Mounting: Crop Compatibility, Height & Row Spacing." PVRack, 27 June 2026. Accessed Aug. 5, 2026. pvrack.com/types/agrivoltaic/.

"Agrivoltaic System Design Tools for Managing Trade-Offs between Energy Production, Crop Productivity and Water Consumption." Environmental Research Letters, IOPscience, 5 Mar. 2024. Accessed June 18, 2026. iopscience.iop.org/article/10.1088/1748-9326/ad2ab8.

"Agrivoltaic System Impacts on Microclimate and Yield of Different Crops within an Organic Crop Rotation in a Temperate Climate." Agronomy for Sustainable Development, Springer Nature Link. Accessed May 24, 2026. link.springer.com/article/10.1007/s13593-021-00714-y.

"Agrivoltaic Systems: Trade-Offs on Microclimate, Physiology, Yield and Canopy Thermal-Spectral Maps." Zenodo, 25 Nov. 2025. Accessed June 7, 2026. zenodo.org/records/18668344.

"Agrivoltaics for Managing Groundwater Overdraft in California: A Suitability Analysis." Environmental Research: Energy, IOPscience, 26 May 2026. Accessed Aug. 14, 2026. iopscience.iop.org/article/10.1088/2976-601X/ae728f.

"Agrivoltaics Revisited: Critical Insights into Shading-Induced Microclimate Change, Yield and Quality, Biodiversity Shifts and Socio-Economic Limitations." AgriEngineering, MDPI, 14 Feb. 2026. Accessed May 30, 2026. www.mdpi.com/2624-7402/8/2/69.

"AGRIVOLTAICS." Global Environment Facility Scientific and Technical Advisory Panel, 17 Jan. 2024. Accessed July 16, 2026. www.thegef.org/sites/default/files/documents/2024-01/EN_GEF.STAP_.C.66.Inf_.04_Agrivoltaics.pdf.

"An Observational Study on the Microclimate and Soil Thermal Regimes under Solar Photovoltaic Arrays." Solar Energy, ScienceDirect, 21 Nov. 2023. Accessed June 25, 2026. www.sciencedirect.com/science/article/abs/pii/S0038092X23007934.

Barron-Gafford, Greg A., et al. "Agrivoltaics Provide Mutual Benefits across the Food-Energy-Water Nexus in Drylands." Nature Sustainability, vol. 2, 2019, pp. 848-855.

"Benefits of Agrivoltaics Across the Food-Energy-Water Nexus." National Renewable Energy Laboratory, 2019. Accessed Aug. 21, 2026. www.nrel.gov/manufacturing/news/program/2019/benefits-of-agrivoltaics-across-the-food-energy-water-nexus.

Bihari, Kinga. "Assessing Agrivoltaic Suitability Across Agricultural Land in California

Through Geospatial Analysis." eScholarship, University of California, 2024. Accessed July 9, 2026. escholarship.org/uc/item/6j4070q4.

Camba, London. "The Effect of Agrivoltaics on Irrigation Water Consumption in Central California." Undergraduate Honors Thesis, Bellarmine University, 2024. Accessed May 6, 2026. scholarworks.bellarmine.edu/ugrad_theses/147/.

Campana, Pietro E., et al. "Scientific Frontiers of Agrivoltaic Cropping Systems." Nature Reviews Clean Technology, 4 Nov. 2025. Accessed Aug. 28, 2026. www.nature.com/articles/s44359-025-00110-9.

"Computational Fluid Dynamics Modelling of Microclimate for a Vertical Agrivoltaic System." Energy Nexus, vol. 9, 2023, article 100173. Accessed June 14, 2026. arxiv.org/abs/2209.01971.

"Cool Tomatoes? Agrivoltaics Could Help California Crop, If the Economics Pan Out." Canary Media, 25 Oct. 2022. Accessed July 27, 2026. www.canarymedia.com/articles/enn/cool-tomatoes-agrivoltaics-could-help-california-crop-if-the-economics-pan-out.

Daniels, Tom, and Iris Zixin Chen. "The Use and Potential of Agrivoltaics in the United States." Kleinman Center for Energy Policy, University of Pennsylvania, 3 Nov. 2025. Accessed May 19, 2026. kleinmanenergy.upenn.edu/research/publications/the-use-and-potential-of-agrivoltaics-in-the-united-states/.

"Drivers of Agrivoltaic Perception in California and North Carolina." Environmental Research: Energy, IOPscience, 8 July 2024. Accessed Aug. 3, 2026. iopscience.iop.org/article/10.1088/2976-601X/ad5449.

"Groundwater in California." Public Policy Institute of California, 10 June 2024. Accessed June 2, 2026. www.ppic.org/publication/groundwater-in-california/.

"Harvesting the Sun Twice: Energy, Food and Water Benefits from Agrivoltaics in East Africa." Renewable and Sustainable Energy Reviews, vol. 208, 2025, article 115066.

Hassanpour Adeh, Elnaz, et al. "Remarkable Agrivoltaic Influence on Soil Moisture, Micrometeorology and Water-Use Efficiency." PLoS ONE, vol. 13, no. 11, 2018, e0203256. Accessed July 22, 2026. www.ncbi.nlm.nih.gov/pmc/articles/PMC6211631/.

"High-Shade Dryland Agrivoltaic Conditions Enhanced Carbon Uptake and Water-Use Efficiency in Zucchini (Cucurbita pepo)." Frontiers in Sustainable Food Systems, 27 May 2026. Accessed June 10, 2026. www.frontiersin.org/journals/sustainable-food-systems/articles/10.3389/fsufs.2025.1686773/full.

"Impact of Agrivoltaics Irrigation and Soil Moisture Dynamics: A Review." ResearchGate, 31 Oct. 2025. Accessed Aug. 10, 2026. www.researchgate.net/publication/396897001.

"Impacts of Agrivoltaic Systems on Microclimate, Water Use Efficiency, and Crop Yield: A Systematic Review." Renewable and Sustainable Energy Reviews, ScienceDirect, 6 June 2025. Accessed June 29, 2026.

www.sciencedirect.com/science/article/abs/pii/S1364032125006033.
"Influence of Agrivoltaic System-Induced Microclimate Modifications on Wheat Growth and Yield." Agricultural and Forest Meteorology, vol. 373, ScienceDirect, 16 Aug. 2025. Accessed July 4, 2026. www.sciencedirect.com/science/article/abs/pii/S0168192325003946.
Lawrence Berkeley National Laboratory. "Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection." Energy Markets & Policy, U.S. Department of Energy, 2024-2026. Accessed May 27, 2026. emp.lbl.gov/queues.
"Lighting the Way for Agrivoltaics: How NREL Empowers Communities to Capture the Benefits of Solar Energy, Agriculture, and Ecosystems." National Renewable Energy Laboratory, 21 Nov. 2024. Accessed Aug. 17, 2026. www.nrel.gov/news/feature/2024/lighting-the-way-for-agrivoltaics.
Liu, Pang-Wei, et al. "Groundwater Depletion in California's Central Valley Accelerates during Megadrought." Nature Communications, vol. 13, 2022, article 7825.
Macy, Abe, et al. "United States Agrivoltaics Policy Landscape: A Review and Research Agenda." Energy Policy, vol. 205, 2025, article 114550.
Marrou, Herve, et al. "Productivity and Radiation Use Efficiency of Lettuces Grown in the Partial Shade of Photovoltaic Panels." European Journal of Agronomy, vol. 44, 2013, pp. 54-66. Accessed June 11, 2026. www.sciencedirect.com/science/article/abs/pii/S1161030112001177.
McCall, James, et al. "Testing Native Pollinator Habitat Seed Mix Establishment at Three Utility-Scale Solar Sites in Minnesota." Environmental Research Communications, National Renewable Energy Laboratory. Accessed July 31, 2026. docs.nrel.gov/docs/fy24osti/88651.pdf.
"More Farmland Going Solar." My Job Depends on Ag Magazine, 30 Apr. 2026. Accessed May 8, 2026. mjdoamag.com/more-farmland-going-solar/.
"New Irrigation Tech Helps Reduce Water Use in Agrivoltaics." pv magazine International, 30 Apr. 2026. Accessed Aug. 24, 2026. www.pv-magazine.com/2026/04/30/new-irrigation-tech-helps-reduce-water-use-in-agrivoltaics/.
Pardo, Miguel A., et al. "Data-Driven Analysis of the Effect of Agrivoltaics Systems on Soil and Air Conditions—A Case Study in Kressbronn, Germany." Applied Sciences, vol. 16, no. 11, 2026, article 5307. Accessed June 20, 2026. www.mdpi.com/2076-3417/16/11/5307.
Sanchez, Geoffrey, et al. "Multi-Year Study of Maize Under Elevated Tracking Agrivoltaic System and Simplified Yield Modeling." npj Sustainable Agriculture, vol. 4, 2026, article 25. Accessed July 13, 2026. www.nature.com/articles/s44264-026-00141-0.
"SGMA Explainer." California Water Impact Network, 24 Mar. 2025. Accessed May 22, 2026. www.c-win.org/cwin-water-blog/2025/3/10/sgma-explainer.
"Soil Suitability and Technical-Construction Criteria for Slope Classification in Agrivoltaic Systems." Solar Energy, ScienceDirect, 7 Feb. 2026. Accessed Aug. 7, 2026.

www.sciencedirect.com/science/article/pii/S0038092X2600023X.
"Solar Energy and Groundwater in the San Joaquin Valley." Public Policy Institute of California. Accessed June 5, 2026. www.ppic.org/publication/solar-energy-and-groundwater-in-the-san-joaquin-valley/.
"Solar Farm Design Guide 2026: Utility-Scale Layout & Optimization." SurgePV, 25 May 2026. Accessed July 25, 2026. www.surgepv.com/blog/solar-farm-design-guide.
Stid, Jacob T., et al. "Impacts of Agrisolar Co-Location on the Food-Energy-Water Nexus and Economic Security." Nature Sustainability, vol. 8, 2025, pp. 702-713. Accessed May 11, 2026. www.nature.com/articles/s41893-025-01546-4.
Stid, Jacob T., et al. "Solar Array Placement, Electricity Generation, and Cropland Displacement across California's Central Valley." Science of the Total Environment, vol. 835, 2022, article 155240.
"The Effect of Agrivoltaic System on Nutrient Content, Yield, and Water Productivity of Potatoes." Frontiers in Horticulture, 16 June 2025. Accessed Aug. 19, 2026. www.frontiersin.org/journals/horticulture/articles/10.3389/fhort.2025.1624013/full.
"The Impact of Land Slope on Agrivoltaics Feasibility." pv magazine Global, 13 Feb. 2026. Accessed June 26, 2026. www.pv-magazine.com/2026/02/13/the-impact-of-land-slope-on-agrivoltaics-feasibility/.
"The Long-Term Suitability of Agrivoltaics as a Climate Adaptation Strategy in the Southwestern United States." Global Environmental Change Advances, vol. 5, 2025, article 100021. Accessed July 8, 2026. www.sciencedirect.com/science/article/pii/S2950138525000099.
"The Techno-Economic Viability of Agrivoltaics for Small-Scale Farmers under Policy and Market Uncertainty." ScienceDirect, 23 June 2026. Accessed July 15, 2026. www.sciencedirect.com/science/article/pii/S2772783126000373.
"The Williamson Act: A Growing Obstacle for Solar and Wind Development in California." Lexology. Accessed Aug. 12, 2026. www.lexology.com/library/detail.aspx?g=165482ec-cf8e-41f1-bbe9-18f49d8a4670.
Trommsdorff, Maximilian, et al. "Cost Comparison between Agrivoltaics and Ground-Mounted PV." pv magazine International, 26 Mar. 2021. Accessed June 17, 2026. www.pv-magazine.com/2021/03/26/cost-comparison-between-agrivoltaics-and-ground-mounted-pv/.
"US Grid Interconnection Backlog Jumps 40%, With Wait Times Expected to Grow as IRA Spurs More Renewables." Utility Dive, 11 Apr. 2023. Accessed July 20, 2026. www.utilitydive.com/news/grid-interconnection-queue-berkeley-lab-lbnl-watt-coalition-wind-solar-renewables/647287/.
"Utility-Scale Solar Fields Can Foster Abundant Biodiversity." National Renewable Energy Laboratory, 2024. Accessed May 4, 2026. www.nrel.gov/manufacturing/news/program/2024/utility-scale-solar-fields-can-foster-abundant-biodiversity.

"Water Evaporation Reduction by the Agrivoltaic Systems Development." AgriSolar Clearinghouse, 27 Oct. 2022. Accessed Aug. 26, 2026. www.agrisolarclearinghouse.org/water-evaporation-reduction-by-the-agrivoltaic-systems-development/.

Widmer, Julian, et al. "Agrivoltaics, a Promising New Tool for Electricity and Food Production: A Systematic Review." Renewable and Sustainable Energy Reviews, vol. 192, 2024, article 114277.

"Williamson Act Contract Removal." California Department of Conservation. Accessed June 9, 2026. www.conservation.ca.gov/dlrp/wa/Pages/removing_contracts.aspx.

"Williamson Act in California: Farmland & Your Rights." Kassouni Law, 11 June 2026. Accessed July 2, 2026. www.kassounilaw.com/2026/06/williamson-act-california/.

Yao, Shiyue, et al. "Optimizing Agrivoltaics Adoption to Support Groundwater Conservation in California's Central Valley." 2025 AAEA & WAEA Joint Annual Meeting, IDEAS/RePEc. Accessed May 17, 2026. ideas.repec.org/p/ags/aaea25/361209.html.

Comparison of Novel Aluminum-Free Deodorants and a Traditional Mineral-Based Deodorant By Joe Fang

Abstract

Personal care items such as deodorants and antiperspirants are applied to the body in order to prevent sweating and eliminate odor. There has been a recent trend toward aluminum-free deodorants, which are popular because consumers want products that contain natural ingredients. There is a lack of scientific investigation into the direct comparison between the performance of aluminum-free and mineral-based antiperspirants in terms of their moisture-blocking capacity. In this study, the stability test, pH test, diffusion barrier test, and simulated-sweat-reduction test were used. The deodorants tested include mineral-based deodorant (Crystal) and two aluminum-free deodorants (Salt & Stone and Rio Deo). In the case of Salt and Stone, the mass increase was minimal (7.89%) compared to that of Rio Deo (13.16%) and Crystal (23.68%). These trends were also reflected in the diffusion barrier test where Crystal had the most dye spread, and Salt and Stone, and Rio Deo were limited in the spread. It was observed in the stability test that the color, texture, smell, and phase separation did not change after two weeks of exposure to room temperature. Overall, the aluminum-free deodorants showed stronger moisture-blocking performance under the experimental conditions. These findings provide preliminary information that may help guide the development and evaluation of alternative deodorant formulations.

Keywords Deodorant; Antiperspirant; Aluminum-free deodorant; Mineral-based deodorant; Sweat reduction; Diffusion barrier; Cosmetic stability; Moisture-control; Plant-base deodorant; Potassium alum

Introduction

Personal deodorants are personal hygiene products designed to eliminate body odor and make people feel more comfortable. The global deodorant market was estimated at approximately $27.8 billion in 2024 and is expected to continue growing (Grand View Research, 2025). Therefore, deodorants are an important category of personal care products, and differences in formulation may influence their performance and suitability for consumers.(Grand View Research, 2025). In general, there are two types of deodorants: antiperspirants, which help prevent sweating, and deodorants, which do not prevent sweating but instead minimize the smell of the body by inhibiting the growth of the odor-causing bacteria.

The most traditional antiperspirants contain aluminum salts, including aluminum chlorohydrate or aluminum zirconium. These chemicals are known to plug sweat glands temporarily and reduce skin perspiration. Despite current scientific evidence and regulatory bodies not having provided any causal evidence between the use of antiperspirants containing aluminum and the occurrence of diseases like breast cancer and Alzheimer disease, consumer trends reflect fears around the exposure to aluminum. As a result,the demand for natural

ingredients and products for sensitive skin has also led to a quick development of aluminum-free deodorants (Teerasumran et al., 2023). Aluminum-free deodorants differ from conventional ones because they usually include plant oil, wax, starch, botanical extracts, probiotics, and other odor-absorbing agents, which help to minimize the odor and enhance the skin comfort without interfering with the body's pores.

In recent years, several studies have focused on deodorant and antiperspirant performance. SweatSENSE was created by Hand et al. (2024) as a paper-based colorimetric sensor that is used to determine the effectiveness of an antiperspirant by measuring its ability to change color when it is wet. This approach demonstrated promise in testing the moisture control of commercial antiperspirants, and this was a fast and cost effective method of testing in the lab. The authors Teerasumran et al. (2023) presented the most recent advances in research on deodorants and antiperspirants, and provided a summary of the common tests performed in this field, such as sweat reduction, antibacterial test, stability test, and skin compatibility test. The review concludes that despite the fact that numerous aluminum-free deodorants made with plant-derived components have been released, there are not many scientific articles that compare the overall quality of these products to traditional mineral-based ones in terms of laboratory conditions. Moreover, Adejokun and Dodou (2020) conducted the experiment on the stability of cosmetic products by storing them at various temperatures and observing the colour changes, texture, pH levels, and phase separation after a period of time. This study showed that the physical stability of a personal care product can be tested in the lab.

Even though these improvements have been made, there is a limited number of studies that have been able to compare the commercial mineral-based and aluminum-free deodorants in direct comparison with each other through various simple laboratory experiments that can be conducted in an educational laboratory. To this end, the stability test, pH testing, diffusion barrier test, and simulated sweat reduction test were used to compare one mineral-based (Crystal) and two aluminum-free (Salt and Stone and Rio Deo) deodorants. The purpose was to make their physical stability and moisture resistance comparable, as well as to determine whether it is possible to detect differences between deodorant formulations using simple laboratory methods.

Materials and Methods

Crystal Mineral Deodorant is a mineral-based deodorant that primarily contains potassium alum. Salt & Stone is an aluminum-free deodorant formulated with plant oils, waxes, probiotics, and botanical extracts. Rio Deo is also aluminum-free and contains tapioca starch, papaya enzyme, and plant-derived oils for odor control and moisture absorption.

The three deodorants were chosen to be tested; the control group was the Crystal Mineral Deodorant, Salt and Stone, and Rio Deo as the experimental groups. Crystal Mineral Deodorant was selected as the mineral-based comparison product because its primary active ingredient is potassium alum. Salt & Stone and Rio Deo were selected as the aluminum-free comparison products because their formulations do not contain aluminum-based antiperspirant salts. Crystal served as the mineral-based comparison group, while Salt & Stone and Rio Deo represented

aluminum-free formulations. This grouping allowed the study to compare moisture-control performance across different formulation approaches.The mineral-based deodorant is known as crystal mineral deodorant, which contains potassium alum as its main ingredient. Salt & Stone is a plant-based deodorant that does not contain aluminum in it. It contains no aluminum, tapioca starch, papaya enzyme, and oils of plants for odor removal and absorption of moisture are also included in Rio Deo.

Four tests were selected to evaluate different characteristics of the deodorants. The stability test assessed whether the formulations changed physically during storage, the pH test characterized their chemical environment, the diffusion barrier test evaluated resistance to liquid spreading, and the simulated sweat reduction test measured moisture uptake under warm and humid conditions. Together, these tests provided complementary information about formulation stability and moisture-control performance.

For the stability test, about 1 g of each deodorant was put in a labeled container and kept at room temperature (25°C) for two weeks. The samples were examined every two days and observations about any change in color, texture, odor, or phase separation were noted.

To test the pH, 10g of each deodorant was added to 90mL of distilled water. The mixtures were placed in a water bath, heated to 50°C for 5-10 minutes, and stirred thoroughly. pH strips were then used to determine and record the pH values of the samples.

In order to assess the performance of a diffusion barrier, a 2 cm circle was marked on filter paper and 0.10 g of deodorant was evenly spread across the indicated area. The treated area was left standing after 5 minutes, and a drop of red ink was placed at the center. The samples were not disturbed for 10 minutes as spreading of dye was observed and documented.

To test the sweat reduction of the simulated sweat, the filter paper was used, and it had a mass of 0.18 g and a mass of deodorant was applied to it at a concentration of 0.20 g resulting in an initial mass of 0.38 g. The samples were put into petri dishes, which were placed on top of warm water (35-37°C) so that they were kept in a humid atmosphere. At the end of 40 minutes, the masses of each sample were recorded. The diameter of dye spreading was used as an indicator of liquid diffusion through or across the treated area. A smaller spreading distance indicated greater resistance to liquid movement under the test conditions.This was determined as the percentage change in the mass and served as the measure of moisture absorption and moisture resistance.



Figure of the 3 antiperspirants

Results

Table 1. Simulated Sweat Reduction Test Results

Product	Initial Mass (g)	Final Mass (g)	Mass Change (g)	Percent Change (%)
Crystal	0.38	0.47	0.09	23.68
Salt & Stone	0.38	0.41	0.03	7.89
Rio Deo	0.38	0.43	0.05	13.16

The simulated sweat reduction test was designed to compare moisture uptake under warm and humid conditions. Table 1 summarizes the results. The original mass of every sample was 0.38 g. Crystal had a final mass of 0.47 g, an increase in mass of 0.09 g (23.68%). Salt & Stone ended up with a mass of 0.41 g, an increase of 0.03 g (7.89%). Rio Deo ended with a mass of 0.43 g and an increase in mass of 0.05 g (13.16%). Overall, Salt and Stone showed the best moisture-blocking ability, followed by Rio Deo and Crystal under the experimental conditions.

Table 2. Diffusion Barrier Test Results

Product	Observation
Crystal	5mm dye spreading
Salt & Stone	3mm dye spreading
Rio Deo	No visible dye spreading observed

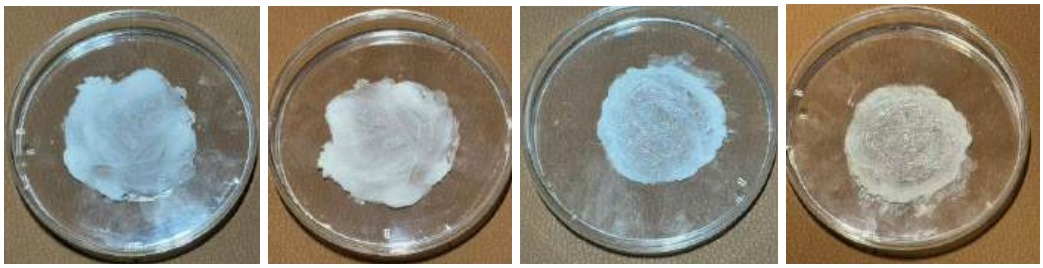
The diffusion barrier test showed differences in liquid spreading among the three products. Crystal had the greatest dye spreading (5 mm), followed by Salt & Stone (3 mm), while Rio Deo showed no visible dye spreading. Overall, the two aluminum-free products showed less dye spreading than Crystal under the experimental conditions.

Table 3. Stability Test Results

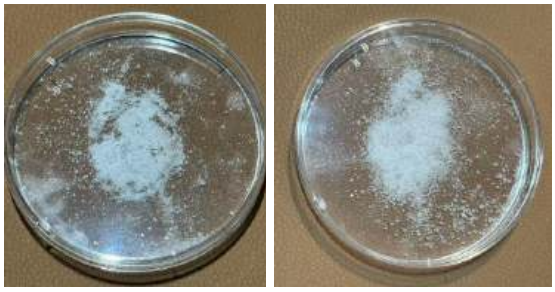
Product	Color Change	Texture Change	Separation
Crystal	None	None	None
Salt & Stone	None	None	None
Rio Deo	None	None	None

Figure 1. Stability test images of Crystal, Salt & Stone, and Rio Deo before testing and after two weeks of storage at room temperature.

(a) Salt & Stone before test (b) Salt & Stone after two weeks (c) Rio Deo before test (d) Rio Deo after two weeks.



(e) Crystal before test (f) Crystal after two weeks



The three deodorants were subjected to a physical stability test, which was conducted by leaving the deodorants in their original appearance after two weeks of being stored at room temperature. The results of the stability tests are presented in Figure 1 and Table 3. There was no change in color, texture, or phase separation of the three deodorants at room temperature over a 2-week period.

Table 4. pH Test Results

Product	pH
Crystal	4.5
Salt & Stone	8.0
Rio Deo	7.5

The pH of all three deodorants was measured using pH strips after sample preparation (Figure 2, Table 4). Crystal had a pH of 4.5, Salt & Stone had a pH of 8.0, and Rio Deo had a pH of 7.5. Overall, Crystal is slightly acidic while Salt & Stone and Rio Deo are slightly basic.

Figure 2. pH test strips and corresponding color changes for Crystal, Salt & Stone, and Rio Deo.

(a) Crystal (b) Salt & Stone (c) Rio Deo (d) pH table





Discussion

The simulated sweat reduction test showed that Salt & Stone had the smallest percentage increase in mass (7.89%), followed by Rio Deo (13.16%) and Crystal (23.68%). The diffusion barrier test showed a similar trend. Crystal allowed the greatest dye spreading (5 mm), Salt & Stone showed limited spreading (3 mm), and Rio Deo showed no visible dye spreading. These results indicate that the three deodorants differed in their ability to reduce moisture movement under the experimental conditions

One possible explanation for the differences observed in moisture uptake and dye spreading is the difference in formulation. The main components of Crystal Mineral Deodorant are mineral salts, whereas Salt and Stone and Rio Deo have plant oils, waxes, and other hydrophobic substances. It has been established that hydrophobic materials slow down water flow through porous surfaces, and this might be the reason why there was a lower mass gain and less dye diffusion in the two aluminum-free deodorants. Other similar paper-based techniques have also been employed by Hand et al. (2024) to measure the performance of antiperspirants showing that colorimetric paper sensors can be used to compare moisture control between products.

Previous studies have identified pH as an important characteristic of deodorant formulations because it can influence skin compatibility, microbial activity, and product stability. (Ali & Yosipovitch, 2013).However, current literature does not provide strong evidence that deodorant pH directly determines sweat reduction or diffusion barrier performance .The study did not find a clear relationship between pH and simulated sweat reduction. Crystal had a pH of 4.5but experienced the biggest increase in mass (23.68%), whereas Salt & Stone had the highest pH (8.) and the smallest increase in mass (7.89%). Rio Deo was similar to Crystal in terms of pH (7.5) and showed an intermediate increase in mass (13.16%). Also, no definite connection was established between pH and diffusion barrier performance. Even though Crystal was the lowest pH deodorant among the three, it exhibited the most dye spreading, and Salt and Stone as well as Rio Deo were slightly higher in their pH values and displayed minimal or no visible dye diffusion.

The diffusion barrier test was found to have a similar pattern to the simulated sweat reduction test. The highest increase in mass and the maximum dye spreading were observed in Crystal, and the least amount of mass and no dye spreading was observed in Salt & Stone. Rio

Deo had neither an observable spread of dyes, nor an intermediate change in mass. These findings indicate that the deodorants that are more effective in moisture blocking were also found to be effective in liquid flow through the filter paper during the experiments. Nevertheless, there is need to conduct further research on this trend using other products as well as repeated trials so as to establish whether the trend can be applied generally.

Importantly, the literature review did not identify peer-reviewed studies that directly compared mineral-based and aluminum-free deodorants using the simulated sweat reduction and diffusion barrier approaches used in this study. Existing research on potassium alum has focused more heavily on antimicrobial activity, formulation, and physical stability (Alzomor et al., 2014; Teerasumran et al., 2023). Therefore, the present study provides preliminary comparative data using simple laboratory methods. Although the results cannot establish which deodorant type is generally superior, they demonstrate a practical approach for identifying performance differences among commercial formulations. The majority of published work on potassium alum has highlighted its antimicrobial activity, formulation, and physical stability, with less emphasis on direct comparisons of moisture-blocking performance (Alzomor et al., 2014; Teerasumran et al., 2023). Thus, this study offers initial comparative data using simple laboratory techniques and underscores the importance of future standardized comparisons between mineral-based and aluminum-free deodorants.

The results of the stability test indicated that none of the deodorants changed in appearance, color, and texture, or phase separation at two weeks of storage. The same findings were made by Adejokun and Dodou (2020), who conducted a cosmetic stability study, which indicated that cosmetics could be physically stable when stored under room temperature conditions. These results indicate that all three products maintained their observable physical characteristics during the two-week room-temperature observation period.

The main significance of this study is that it demonstrates how simple laboratory tests can reveal measurable differences between commercial deodorant formulations. The stronger moisture-blocking performance observed in Salt & Stone and Rio Deo suggests that formulation characteristics, rather than product category alone, may be important when developing aluminum-free deodorants. For future product development, characteristics such as hydrophobic ingredients, resistance to liquid spreading, physical stability, and skin-compatible pH may all need to be considered together rather than focusing on a single property. These findings may also help consumers understand that “aluminum-free” and “mineral-based” describe different formulation approaches and do not by themselves indicate how well a product will control moisture. However, the preliminary nature of this study means that these implications should be tested with more products and repeated experiments.

The limitations of the study are as follows. First, three commercial deodorants were tested, and each was a single formulation. Second, filter paper was used in place of human skin in the simulated sweat reduction test; the experimental conditions did not allow for the complete replication of the real sweating process. Moreover, only one trial was conducted for each product, which might undermine the credibility of the findings. The future research may involve

more brands of deodorants, multiple experiments, artificial sweat solutions, and quantitative image analysis to gain more detailed and accurate information.

Conclusion

This research compared the efficacy of a single mineral-based deodorant (Crystal) and two non-aluminum-based deodorants (Salt and Stone and Rio Deo). The study aimed to compare the physical stability and their ability to block moisture using simple experiments. There were four experiments, namely, the test of stability, pH test, the diffusion barrier test, and the sweat reduction simulation test.

It was determined that the three deodorants did not undergo any physical changes and there were no visible alterations of color, texture, smell or phase change in them after two weeks of storage. The pH levels of the products varied from acidic to slightly alkaline. Crystal had the most dye diffusion during the diffusion barrier test, with the least spreading of Salt & Stone and Rio Deo. The percentage mass growth of the product was the least in Salt and stone as compared to other products, such as Rio Deo and Crystal, in the simulated sweat reduction experiment.

There was no obvious connection between the pH and deodorant performance. Nevertheless, there was a consistent pattern between the simulated sweat reduction test and diffusion barrier test, which indicates that products with more effective moisture-blocking performance also tended to reduce liquid diffusion through the filter paper.

In general, under the experimental conditions applied in this experiment, the two aluminum-free deodorants proved to be more effective in terms of moisture-blocking than the mineral-based deodorant, whereas all three products were physically stable. The procedures adopted in this project provided a cheap and easy way to compare deodorant performance.

More broadly, the findings suggest that deodorant performance may depend on the specific formulation and its physical properties rather than simply whether a product is aluminum-free or mineral-based. The study also demonstrates that relatively simple laboratory methods can be used to screen differences in commercial deodorant formulations. These findings may be useful for preliminary product development and for helping consumers understand that ingredient category alone does not determine moisture-control performance. Because this study used only three products and one trial per product, further research with more formulations, repeated trials, artificial sweat, and quantitative measurements is needed to confirm these findings.

Acknowledgement

I would like to acknowledge my mentor, Nathanial, for his guidance and helpful advice on this research project, including experimental planning, data interpretation, and scientific writing. I am grateful to my school lab for supplying the necessary equipment and supplies to conduct this study. I would also like to thank my family and friends for their encouragement and support during the research process.

Works Cited

Grand View Research. (2025). Deodorants market size, share & trends analysis report. https://www.grandviewresearch.com/industry-analysis/deodorants-market

Teerasumran P, Velliou E, Bai S, Cai Q. Deodorants and antiperspirants: novel trends in activities and testing methods. Int J Cosmet Sci. 2023;45(4):426443. https://doi.org/10.1111/ics.12852

Hand, R. A., Efstathiou, S., Wemyss, A. M., Grypioti, M., Kirby, G., Barlow, T., Tinley, E. C., Ford, J., Jamieson, A., Reynolds, J., Miller, J., Bates, S., Khoshdel, E., and Haddleton, D. M. (2024). The easiest way of assessing the effectiveness of antiperspirants on paper-based colorimetric paper sensors: SweatSENSE. https://doi.org/10.1039/D3LP00214D.

Ali, S. M., & Yosipovitch, G. (2013). Skin pH: From basic science to basic skin care. Acta Dermato-Venereologica, 93(3), 261–267. https://doi.org/10.2340/00015555-1531

Adejokun, D. A., & Dodou, K. (2020). New approach to the long-term stability of cream formulations with natural oils. Cosmetics, 7(4), 86. https://doi.org/10.3390/cosmetics7040086

International Conference on Harmonisation of Technical Requirements for Registration of Pharmaceuticals for Human Use. (2003). ICH harmonized tripartite guideline Q1A(R2): Stability testing of new drug substances and products. https://database.ich.org/sites/default/files/Q1A%2528R2%2529%20Guideline.pdf

Investigating the Efficiency of Melanin as an Active Layer Absorber in Perovskite-Inspired Solar Cell Architecture via SCAPS-1D Simulation By Hamza Ahmed

Abstract

Melanin is a broadband biological pigment with exceptional photostability and sustainable sourcing, yet its potential as a solar cell absorber remains largely unexplored due to its characteristically high recombination rates and poor charge mobility. This study presents the first computational investigation to exceed 4% power conversion efficiency (PCE) using melanin as the active layer absorber, a substantial improvement over previously reported melanin-based and bio-sensitized photovoltaic devices. Using SCAPS-1D numerical simulation under standard test conditions, a perovskite-inspired p-i-n device architecture (Left Contact/Cu_2O/Eumelanin/TiO_2/Right Contact) was employed to overcome melanin's recombination limitations by facilitating sub-picosecond charge separation at the transport layer interfaces. Sequential optimisation of the hole transport layer, electron transport layer, and melanin absorber thicknesses yielded optimal values of 0.550 μm, 0.020 μm, and 0.050 μm respectively, producing an optimised device with a Voc of 1.0941 V, Jsc of 9.0162 $mAcm^{-2}$, Fill Factor of 48.69% and PCE of 4.80%. Sensitivity analyses revealed that carrier mobility and interface defect density are the dominant performance-limiting parameters, highlighting the necessity of fabricating high-quality, low-defect melanin films to achieve viable device performance. These findings establish a clear design framework for future experimental realisation of melanin-based photovoltaics and position bio-derived absorbers as viable candidates for next-generation sustainable solar energy conversion.

Introduction

The global photovoltaic market is dominated by silicon-based solar cells, which currently account for over 90% of worldwide solar energy production. However, the energy-intensive manufacturing process of silicon presents a significant sustainability challenge: the production of 1 kg of metallurgical grade silicon requires 14–16 kWh of energy (Hoffs). As global energy demand continues to rise, the development of low-cost, sustainable alternatives to silicon has become increasingly urgent. Organic photovoltaics (OPVs) have emerged as a promising candidate in this regard, offering key advantages including low-cost fabrication, mechanical flexibility, biocompatibility, and significantly reduced energy consumption during manufacturing compared to conventional inorganic devices.

Among the various approaches explored within OPVs, the use of naturally derived biological pigments as light-harvesting sensitizers has attracted considerable interest. These bio-derived dyes, which include anthocyanins, chlorophyll derivatives, and carotenoids, mimic natural photosynthesis by absorbing incident photons and generating electron-hole pairs that drive an external photocurrent. Despite their sustainability advantages, bio-sensitized devices have consistently demonstrated poor power conversion efficiencies (PCEs), largely due to narrow absorption spectra, rapid dye degradation, and high recombination rates.

Anthocyanin-based dye-sensitized solar cells (DSSCs), for example, have reported efficiencies as low as 0.009%, with values rarely exceeding 1% even when combined with copigments (Pramananda et al.). These fundamental limitations have motivated the search for bio-derived absorbers with broader spectral response and greater intrinsic stability.

Melanin is a ubiquitous biological pigment found in the skin, hair, eyes, and feathers of most animals, where it serves as a natural photoprotectant. Unlike narrow-band biological dyes, melanin exhibits a broadband absorption spectrum spanning 200 nm to 1100 nm, covering a substantial portion of the solar spectrum (Daugman). Its amorphous, disordered molecular structure, comprising eumelanin monomers DHI and DHICA, gives rise to an extensive system of conjugated π bonds that enables efficient photon capture across multiple wavelengths. Furthermore, melanin demonstrates exceptional long-term photostability, making it an attractive candidate for sustainable photovoltaic applications. Despite these favourable properties, melanin has remained largely underexplored as a solar cell absorber, primarily due to its characteristically poor charge mobility and high non-radiative recombination rates, which have historically limited device performance to PCEs of around 0.10% (Al-Shamery et al.).

The central challenge in melanin-based photovoltaics is therefore the management of photogenerated charge carriers, specifically preventing their recombination before they can be collected at the electrodes. Perovskite solar cells have successfully addressed analogous charge management challenges through the use of dedicated hole transport layers (HTLs) and electron transport layers (ETLs) that facilitate rapid charge separation and selective carrier extraction. Inspired by this architecture, the present study proposes a p-i-n device configuration in which melanin serves as the intrinsic absorber layer, sandwiched between Cu_2O as the HTL and TiO_2 as the ETL. This configuration is designed to exploit the favourable band alignment between these materials to achieve sub-picosecond charge separation, effectively suppressing the recombination losses that have previously limited melanin-based devices.

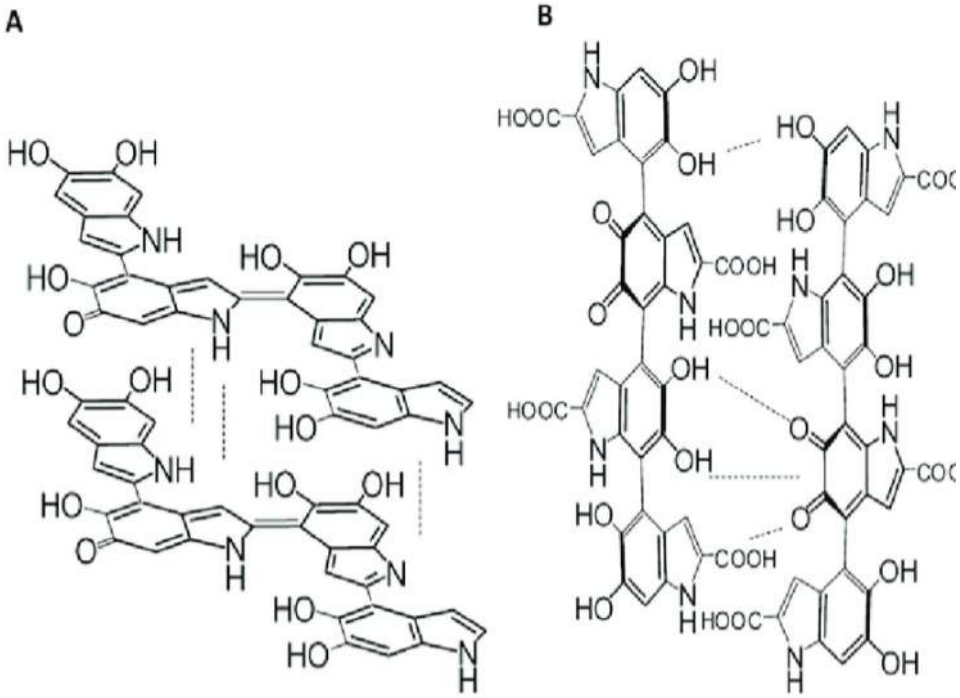
To investigate the viability of this architecture, SCAPS-1D numerical simulation was employed, a widely used and validated tool for the computational modelling of thin film solar cells (Niemegeers et al.). Through sequential optimisation of layer thicknesses and sensitivity analyses of key material parameters, this study aims to establish the performance potential of melanin as a photoabsorber and identify the primary material challenges that must be addressed for future experimental realisation. To our knowledge, this is the first computational study to exceed 4% PCE using melanin as the active layer absorber, representing a significant advancement in the field of bio-derived photovoltaics.

Melanin as an Active Layer Absorber

Melanin is an ubiquitous biological pigment responsible for the dark brown and black coloration in most animals, plants, and fungi. In nature, melanin is primarily found in the skin, hair, feathers, and the iris of the eye, where it serves as a primary photoprotectant. Most natural melanin is categorized as eumelanin, which is derived from two main monomeric units:

5,6-dihydroxyindole (DHI) and 5,6-dihydroxyindole-2-carboxylic acid (DHICA). The amorphous, disordered state of melanin contributes to its unique photophysical properties.

Melanin has a broad range of absorption spectrum from 200 nm to 1100 nm, with a peak at around 335 nm. The extensive system of conjugated π bonds allows for the efficient capture of photons across the solar spectrum. However, melanin typically does suffer from high recombination rates due to its disordered nature.



(A) DHI melanin represented as planar oligomeric scaffolds, (B) DHICA melanin is made up of twisted linear oligomer structures featuring atropisomerism.

Figure 1. Supramolecular structures of DHI and DHICA melanin (Martinez et al.)

Methods

In this study, I employed the advanced SCAPS-1D software (Solar Cell Capacitance Simulator in One Dimension) to conduct rigorous simulations investigating the usefulness of eumelanin as an active layer absorber. SCAPS-1D was developed at the Department of Electronics and Information Systems (ELIS) of the University of Gent, Belgium. The simulations were conducted under Standard Test Conditions (STC) at 300 K, a solar irradiance of 1000 W/m², and Air Mass 1.5G. The primary objective was to find an optimal configuration using melanin as photoabsorber. The simulated device followed a p-i-n configuration: Left Contact/HTL/Melanin/ETL/Right Contact. HTL stands for Hole Transport Layer and ETL stands for Electron Transport Layer.

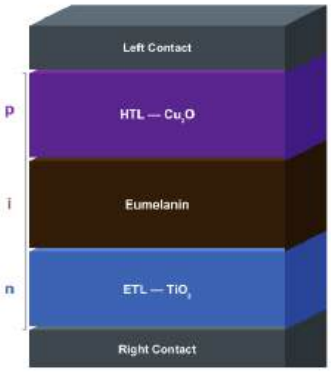


Figure 2. Schematic diagram of the p-i-n solar cell configuration (Left Contact/Cu_2O/Eumelanin/TiO_2/Right Contact)

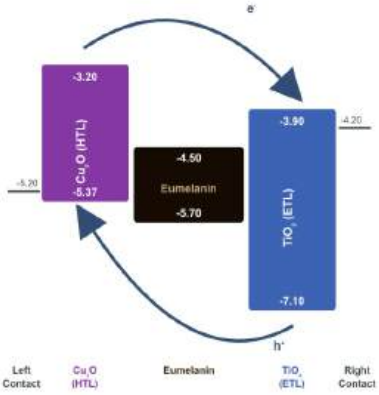


Figure 3. Energy band diagram of the simulated p-i-n device architecture

Baseline Configuration

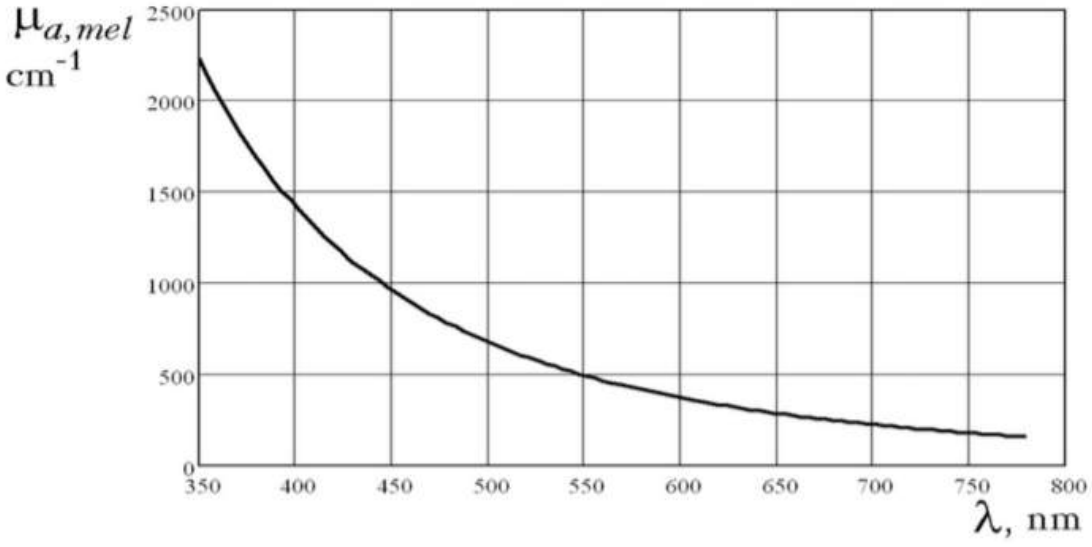


Figure 4. The absorption coefficient of melanin plotted against the wavelength of light (Petruk et al.). Data from this graph was used to model the absorption spectrum of melanin in the simulation.

The parameters of the metal contacts are shown in the table below. Flat band conditions were assumed at both contacts, implying ideal ohmic behaviour with no interfacial barrier between the transport layers and the electrodes

	Surface recombination velocity for electrons (cm/s)	Surface recombination velocity for holes (cm/s)
Left Contact	1.00E+3	1.00E+5
Right Contact	1.00E+3	1.00E+5

Table 1. Parameters for the left and right contacts

Simulations began with an initial baseline configuration based on material parameters sourced from published literature (Altindemir and Gumus; Meredith and Sarna; Etafa Tasisa et al.; Sun et al.; Chauhan and Patel; Obukhova et al.; Ranjan et al.; Rahman; Krasienapibal et al.; Sekkat et al.). Each layer had a type neutral defect of 1×10^{14} cm^{-2}. As this study is entirely computational, experimental characterisation requirements such as EQE measurements, device calibration, and error statistics from physical devices are not applicable.

Layer	Bandgap (eV)	Electron affinity (eV)	Dielectric permittivity (relative)	CB effective DOS ($1/cm^3$)	VB effective DOS ($1/cm^3$)	Electron mobility (cm^2/Vs)	Hole mobility (cm^2/Vs)	ND ($1/cm^3$)	NA ($1/cm^3$)
Cu_2O (HTL)	2.17	3.2	7.11	2.02E+17	1.10E+19	2.00E+02	8.00E+01	0	1.00E+18
Melanin	1.2	4.5	10.0	1.00E+19	1.00E+19	1.00E-04	1.00E-04	0	1.00E+17
TiO_2 (ETL)	3.2	3.9	9.00	1.00E+21	2.00E+20	2.00E+01	1.00E+01	2.00E+19	0

Table 2. Parameters of the simulated device layers

Sequential Layer Thickness Optimisation and Sensitivity Analysis

A sequential optimisation procedure was employed to identify the optimal thickness for each layer. In the first step, the HTL thickness was varied systematically from 0.100 μm to 0.550 μm while the melanin and ETL thicknesses were held at their baseline values. The Voc, Jsc, FF and PCE were recorded at each thickness and the value yielding the highest PCE was identified as the optimal HTL thickness, which was then fixed for all subsequent simulations.

In the second step, the ETL thickness was varied from 0.020 μm to 0.200 μm while the HTL was fixed at its optimised value and the melanin thickness was held at the baseline value. The optimal ETL thickness was identified using the same criterion and fixed for all subsequent simulations.

In the third step, the melanin absorber thickness was varied from 0.040 μm to 0.070 μm while both the HTL and ETL were fixed at their respective optimised values. The optimal

melanin thickness was identified and adopted as the final absorber thickness for the fully optimised device configuration.

Following thickness optimisation, two sensitivity analyses were conducted using the fully optimised device configuration. First, the melanin carrier mobility, comprising both electron and hole mobility, maintained equal throughout, was varied from 10^{-4} cm²/Vs to 10^{-7} cm²/Vs to assess the sensitivity of device performance to charge transport properties of the absorber. Second, the interface defect density at both the HTL/melanin and melanin/ETL interfaces was varied from 10^{10} cm^{-2} to 10^{14} cm^{-2}, with all other parameters held constant, to evaluate the impact of interfacial recombination on device performance.

Results & Discussion

Effect of varying HTL thickness on PCE

The thickness of the HTL layer was systematically varied from 0.100 μm to 0.550 μm, while the melanin and ETL thickness remained constant at the baseline values. For each thickness, the Voc, Jsc, FF and PCE were recorded.

HTL/μm	Voc/V	Jsc/mAcm^{-2}	FF/%	PCE/%
0.100	0.9618	3.558	53.66	1.84
0.150	1.0038	4.780	51.91	2.49
0.200	1.0308	5.7547	50.87	3.02
0.300	1.0626	7.1695	49.74	3.79
0.350	1.0725	7.6849	49.40	4.07
0.400	1.0801	8.1079	49.15	4.30
0.450	1.0860	8.4578	48.96	4.50
0.500	1.0907	8.7492	48.81	4.66
0.550	1.0495	8.9938	48.69	4.79

Table 3. Effect of varying Cu_2O (HTL) thickness on photovoltaic performance parameters of the simulated device.

Beyond 0.550μm the increase in PCE was marginal. The most significant effect of increasing HTL thickness was observed in Jsc, increasing by 152.8% as HTL varies from 0.100μm to 0.550μm. As the HTL thickness increases, the layer serves as a more effective barrier for electrons, reducing the probability of photogenerated holes recombining with electrons. As a result, the Jsc increases. On the other hand, FF decreases by 9.26% when HTL thickness is increased, since a greater thickness leads to a higher series resistance and thus the holes have a greater distance to travel before reaching the contact, leading to increasing ohmic

losses. The Voc experiences a modest increase, owing to the reduced recombination rate. 0.550μm was identified as the ideal HTL thickness and used for further simulations that investigate the effect of varying ETL and melanin thickness.

Effect of varying ETL thickness on PCE

The thickness of ETL was varied from 0.150μm to 0.060μm while the thickness of HTL was kept as 0.550μm and thickness of the melanin layer was kept as 0.050μm. All other parameters were the same as the baseline configuration.

ETL thickness/μm	Voc/V	Jsc/mAcm^{-2}	FF/%	PCE/%
0.200	1.0957	8.8215	48.73	4.71
0.150	1.0956	8.8703	48.72	4.73
0.100	1.0953	8.9241	48.70	4.76
0.080	1.0951	8.9469	48.70	4.77
0.060	1.0948	8.9703	48.69	4.78
0.040	1.0945	8.9938	48.69	4.79
0.020	1.0941	9.0162	48.69	4.80

Table 4. Effect of varying TiO_2 (ETL) thickness on photovoltaic performance parameters of the simulated device.

As opposed to the HTL, the ETL thickness does not have as significant of an effect on the PCE. As the ETL thickness decreases from 0.200μm to 0.020μm, the PCE experiences an increase of 1.91%. 0.020μm was chosen as the optimum thickness for ETL.

Effect of varying melanin thickness on PCE

The thickness of melanin was systematically varied from 0.050μm to 0.010μm while the thickness of HTL was kept at 0.550μm and the thickness of ETL was kept at 0.020μm. All other parameters were the same as those in the baseline configuration.

Melanin thickness/μm	Voc/V	Jsc/mAcm^{-2}	FF/%	PCE/%
0.070	0.9257	8.3792	41.86	3.25
0.060	0.9886	8.7458	46.40	4.01
0.050	1.0941	9.0162	48.69	4.80
0.040	1.2887	9.2372	47.04	5.60

Table 5. Effect of varying melanin thickness on photovoltaic performance parameters of the simulated device.

Thickness including and below 0.040μm produced Voc values greater than 1.2V, which is greater than the bandgap of melanin and thus, physically impossible. The Jsc experienced an

increase of 7.72% as the melanin thickness decreased from 0.070μm to 0.050μm. Contrary to ordinary behavior, Jsc increases with decreasing absorber layer thickness. This result is attributed to the low carrier mobility of melanin. At lower thicknesses, the electron-hole pairs have to travel a shorter distance to reach the interface, which reduces the recombination rate such that it outweighs the lower photon absorption by the thinner absorber layer. 0.050μm was identified as the optimum thickness.

Effect of varying melanin mobility on PCE

The simulation was carried out at a melanin thickness of 0.050μm, ETL thickness of 0.020μm and a HTL thickness of 0.550μm. Melanin carrier mobility i.e. the electron mobility and hole mobility was varied between 10^{-4} cm²/Vs to 10^{-7} cm²/Vs and the results were recorded.

Melanin carrier mobility/cm²/Vs	Voc/V	Jsc/mAcm^{-2}	FF/%	PCE/%
10^{-4}	1.0941	9.0162	48.69	4.80
10^{-5}	0.8367	7.0471	29.99	1.77
10^{-6}	0.5447	1.9282	16.16	0.17
10^{-7}	0.4595	135.1497E-3	6.40	0.004

Table 6. Effect of varying melanin carrier mobility on photovoltaic performance parameters of the simulated device.

The results confirm that carrier mobility has a very significant effect on device performance. The observed degradation in all photovoltaic parameters with decreasing carrier mobility is due to increase in non-radiative recombination rates within the melanin absorber layer.

Effect of varying interface defect density on PCE

The interface defect density was varied from 10^{10} cm^{-2} to 10^{14} cm^{-2} while the melanin thickness was kept at 0.050μm, HTL thickness was kept at 0.550μm and ETL thickness was kept at 0.020μm.

Defect Density/cm^{-2}	Voc/V	Jsc/mAcm^{-2}	FF/%	PCE/%
10^{10}	1.0941	9.0162	48.69	4.80
10^{11}	1.0919	8.9730	48.60	4.76
10^{12}	1.0725	8.5627	47.81	4.39
10^{13}	0.9718	5.8735	43.52	2.48
10^{14}	0.8098	1.4120	40.61	0.46

Table 7. Effect of varying interface defect density photovoltaic performance parameters of the simulated device.

To contextualize the performance of the optimised device, Table 8 compares the efficiency of this device with previously reported bio-sensitized devices (Mejica et al.).

Absorber	PCE / %
Eumelanin (this work)	4.80
Perilla (chlorophyll)	0.50
Sesbania grandiflora (anthocyanin)	1.02
Achiote (carotenoid)	0.37
Punica granatum peel (flavonoid)	1.86

Table 8. Comparing the efficiency of the optimised device with previously reported bio-sensitized devices.

This behaviour can be attributed to trap-assisted Shockley-Read-Hall (SRH) non-radiative recombination as the dominant loss mechanism. At low defect densities, trap states are sufficiently sparse that photogenerated carriers can navigate the absorber and reach the transport layer interfaces without significant recombination losses. As defect density increases beyond 10^{12} cm^{-2}, the concentration of trap states within the melanin bandgap becomes large enough to substantially increase the probability of carrier capture and non-radiative recombination, reducing the quasi-Fermi level splitting and carrier collection efficiency simultaneously.

The optimised p-i-n device architecture yielded a Voc of 1.0941 V, Jsc of 9.0162 $mAcm^{-2}$, Fill Factor of 48.69%, and a PCE of 4.80%. The relatively high simulated PCE is attributable to the perovskite-inspired device configuration, which employs dedicated hole and electron transport layers before and after the melanin absorber respectively. This architecture directly addresses melanin's primary limitation (high recombination rates) by facilitating rapid charge separation at the Cu_2O/melanin and melanin/TiO_2 interfaces before photogenerated carriers can undergo non-radiative relaxation. The fill factor of 48.69% is on the lower end, attributable to melanin's poor intrinsic charge mobility which increases series resistance and recombination losses within the absorber layer. The modest Jsc of 9.0162 $mAcm^{-2}$ reflects the constraint imposed by the thin optimal melanin absorber thickness of 0.050 μm; while a thicker absorber would increase photon absorption. It demonstrates that increased melanin thickness leads to greater recombination losses that outweigh the benefit of additional light absorption, due to the low carrier mobility.

Moreover, carrier mobility and interface defect density are the dominant performance-limiting parameters. Reducing mobility from 10^{-4} to 10^{-5} cm^2/Vs causes PCE to collapse from 4.80% to 1.77%, while interface defect densities exceeding 10^{12} cm^{-2} cause PCE to fall by over 90%. These results directly explain the discrepancy between the simulated PCE and experimentally reported values of approximately 0.10%. As shown in Table 8, the optimised device nevertheless surpasses all representative bio-sensitized devices reported in literature,

demonstrating the significant potential of the perovskite-inspired p-i-n architecture for bio-derived photovoltaics.

Before contextualising these results further, the limitations of this study must be acknowledged. The carrier mobility of melanin was set at 10^{-4} cm²/Vs, corresponding to the upper bound of hydrated melanin. In practice, mobility is highly sensitive to hydration state, with dehydrated melanin exhibiting values as low as 10^{-8} cm²/Vs, and hydrated melanin ranging from 10^{-4} to 10^{-6} cm²/Vs. Additionally, the absorption coefficient data used to model the melanin absorber layer was derived from optical measurements of melanin within human skin tissue, which may differ from laboratory-prepared eumelanin films due to differences in melanin concentration, hydration state, surrounding matrix, and degree of aggregation. This introduces uncertainty into the simulated photocurrent generation and may cause the modelled absorption to deviate from that of an isolated eumelanin thin film. Furthermore, this study assumes ideal ohmic electrode contacts through flat band conditions, implying no interfacial energy barrier between the transport layers and the electrodes. In practice, real metal-semiconductor interfaces exhibit non-ideal behaviour including Schottky barriers, contact resistance, and interface recombination, all of which would reduce Voc, FF and overall PCE. Collectively, these idealised assumptions suggest that the simulated PCE of 4.80% represents a theoretical upper bound, and experimental realisation of this architecture would likely yield significantly lower efficiencies until advances in melanin film processing, interface engineering, and contact optimisation are achieved.

Several directions exist for future research building on this work. The most immediate priority is the experimental fabrication of the proposed Cu_2O/Eumelanin/TiO_2 p-i-n architecture, which would validate the computational findings presented here and quantify the gap between simulated and real device performance. On the materials side, the development of cost-efficient synthetic analogues of melanin with higher charge mobility and lower series resistance represents a promising avenue. The incorporation of light-trapping textures into the device architecture should also be explored, as these would allow a thinner melanin absorber layer to capture more photons without the series resistance penalty associated with increased absorber thickness, directly addressing the FF-Jsc trade-off identified in this study. Future computational studies should additionally investigate alternative HTL and ETL materials with more favourable band alignment to melanin.

Conclusion

This study presents the first computational investigation to exceed 4% power conversion efficiency using melanin as the active layer absorber in a solar cell, achieving a PCE of 4.80% with a Voc of 1.0941 V, Jsc of 9.0162 mAcm^{-2}, and Fill Factor of 48.69% through sequential optimisation of layer thicknesses in a perovskite-inspired p-i-n architecture. The results demonstrate that the primary challenge of melanin-based photovoltaics (high recombination rates) can be significantly mitigated through the incorporation of dedicated Cu_2O hole transport and TiO_2 electron transport layers, which facilitate rapid charge separation at the absorber interfaces. Sensitivity analyses establish that carrier mobility and interface defect density are the

dominant performance-limiting parameters, with PCE collapsing by over 99% as mobility decreases from 10^{-4} to 10^{-7} cm^2/Vs and by over 90% as interface defect density exceeds 10^{12} cm^{-2}, providing concrete targets for future experimental realisation. While the simulated PCE represents a theoretical upper bound under idealised conditions, it nonetheless surpasses all representative bio-sensitized devices reported in literature and represents a substantial advancement over previously reported experimental melanin-based devices. The advantages of melanin, including its broadband absorption, exceptional photostability, biocompatibility, and sustainable sourcing, position it as a compelling candidate for next-generation bio-derived photovoltaics. This study therefore serves as a design framework and benchmark for future experimental and computational work aimed at realising the full potential of melanin as a sustainable solar energy converter.

Works Cited

Al-Shamery, Noah, et al. "From black pigment to green energy: shedding light on melanin electrochemistry in dye-sensitized solar cells." Materials Advances, vol. 6, no. 10, 2025, pp. 3073-3083. doi:10.1039/d5ma00081e.

Altindemir, G., and C. Gumus. "Cu2O thin films prepared by using four different copper salts at a low temperature: An investigation of their physical properties." Materials Science in Semiconductor Processing, vol. 107, 2020, pp. 104805. doi:10.1016/j.mssp.2019.104805.

Chauhan, Khushbu R., and Dipal B. Patel. "Functional nanocrystalline TiO2 thin films for UV enhanced highly responsive silicon photodetectors." Journal of Alloys and Compounds, vol. 792, 2019, pp. 968-975. doi:10.1016/j.jallcom.2019.04.111.

Daugman, John. "Absorption Spectrum of Melanin." University of Cambridge Computer Laboratory, www.cl.cam.ac.uk/~jgd1000/melanin.html.

Etafa Tasisa, Yonas, et al. "Band gap engineering of titanium dioxide (TiO2) nanoparticles prepared via green route and its visible light driven for environmental remediation." Results in Chemistry, vol. 11, 2024, pp. 101850. doi:10.1016/j.rechem.2024.101850.

Glennise Faye C. Mejica, , et al. "Natural dye (chlorophyll, anthocyanin, carotenoid, flavonoid) photosensitizer for dye-sensitized solar cell: A review." Maejo International Journal of Energy and Environmental Communication, vol. 4, no. 1, 2022, pp. 12-22. doi:10.54279/mijeec.v4i1.247970.

Hoffs, Charlie. "How Are Solar Panels Made?" Union of Concerned Scientists Blog, www.blog.ucs.org/charlie-hoffs/how-are-solar-panels-made/.

Krasienapibal, Thantip S., et al. "Improved room temperature electron mobility in self-buffered anatase TiO2 epitaxial thin film grown at low temperature." Japanese Journal of Applied Physics, vol. 53, no. 9, 2014, pp. 090305. doi:10.7567/JJAP.53.090305.

Martínez, Luz María, et al. "Biotechnological Production of Melanins with Recombinant Microorganisms." Melanins: Functions, Biotechnological Production, and Applications, Springer International Publishing, 2023, pp. 67-90. doi:10.1007/978-3-031-27799-3_4.

Meredith, Paul, and Tadeusz Sarna. "The physical and chemical properties of eumelanin." Pigment Cell Research, vol. 19, no. 6, 2006, pp. 572-594. doi:10.1111/j.1600-0749.2006.00345.x.

Niemegeers, Alex, et al. SCAPS-1D. Version 3.8, University of Ghent, 2020, scaps.elis.ugent.be/.

Obukhova, T, et al. "Optical studies of melanin films as a material for solar light absorbers." Materials Research Express, vol. 12, no. 6, 2025, pp. 065503. doi:10.1088/2053-1591/ade227.

Petruk, Vasyl, et al. "The optical diagnostics of parameters of biological tissues of human intact skin in near-infrared range." SPIE Proceedings, vol. 10031, 2016, pp. 100313C. doi:10.1117/12.2249345.

Pramananda, Viqry, et al. "Anthocyanin as natural dye in DSSC fabrication: A review." IOP Conference Series: Materials Science and Engineering, vol. 1122, no. 1, 2021, p. 012104. doi:10.1088/1757-899X/1122/1/012104.

Rahman, M. Atowar. "Design and simulation of a high-performance Cd-free Cu2SnSe3 solar cells with SnS electron-blocking hole transport layer and TiO2 electron transport layer by SCAPS-1D." SN Applied Sciences, vol. 3, no. 2, 2021. doi:10.1007/s42452-021-04267-3.

Ranjan, Rahutosh, et al. "SCAPS study on the effect of various hole transport layers on highly efficient 31.86% eco-friendly CZTS based solar cells." Scientific Reports, vol. 13, no. 1, 2023. doi:10.1038/s41598-023-44845-6.

Sekkat, Abderrahime, et al. "Open-air printing of Cu2O thin films with high hole mobility for semitransparent solar harvesters." Communications Materials, vol. 2, no. 1, 2021. doi:10.1038/s43246-021-00181-8.

Sun, Bin, et al. "Numerical investigation of the Cu2O solar cell with double electron transport layers and a hole transport layer." Optical Materials, vol. 131, 2022, pp. 112642. doi:10.1016/j.optmat.2022.112642.

Predicting Weather Patterns for Rochester, New York By Wohaan Dissanayake

Introduction

Climate change is one of the most important environmental problems. Multiple studies show that climate patterns are shifting, leading to global warming, and that extreme weather patterns are damaging agriculture and human habitats. So, it is important to study climate and develop predictive models to understand future climate trends and plan accordingly. There are few studies that analyze long-term climate data and develop predictive models at the local level.

Rochester, New York, is located on the southern shore of Lake Ontario. Studying the Rochester climate will represent the impact of climate change on the Great Lakes basin. The city's weather is characterized by seasonal variations, lake-effect snowfall, and a humid continental climate. Studying long-term climate trends in this region is important for agriculture practices, water resource management, urban planning, and public health preparedness. Studies show that the Great Lakes region has experienced approximately 2.3°F of accelerated warming since 1951, exceeding the global average[1]. Rochester has had a 1.8°F increase in annual mean temperature since 1970, with warming winter months[2]. The Northeastern United States has shown measurable warming and changes in snowfall. There are specific weather patterns noted in cities in the Great Lakes region, including Rochester, New York, known as the lake-effect weather pattern, as these cities are located in close proximity to large bodies of water. In this study, we analyzed weather trends of Rochester, New York, over a century and created a predictive model to forecast temperature, precipitation, and snowfall patterns throughout the year.

Research Objectives

1. Develop a predictive model for Rochester weather patterns (2024-2125)
2. Created climate predictions for several emission pathways.
3. Validate model predictions via recent recorded data.
4. Provide uncertainty measurement through confidence intervals.

Climate prediction models advanced from simple linear extrapolations into complex earth system models. Early models relied on statistical trend analysis[3],while newer models include physical processes, feedback mechanisms, and socioeconomic scenarios[4]. For regional predictions, downscaling techniques bridge the gap between climate models and local impact[5].

In previous studies of the Great Lake region, multiple key climate trends were identified, such as accelerated warming relative to the global average[6], reduced ice cover on the lakes[7], changing precipitation trends with more extreme events[8], and changes in lake-effect snowfall patterns[9].

Rochester climate research shows a 15% reduction in annual snowfall since 1970[10], more frequent heat waves[11], and shifting growing seasons[12]. There are numerous studies on global and

regional climate change, but there are no century-scale projections for specific cities such as Rochester. This predictive model fills this gap.

Methodology

Data

We used historical weather data for Rochester, 1900-2023, which was collected from a publicly available dataset maintained by the National Oceanic and Atmospheric Administration (NOAA), which includes:

Precipitation measurements
Daily Temperature (minimum, maximum and average)
Snowfall
Extreme weather event frequencies

Data Quality control procedures include missing or trace (for precipitation) values in the dataset are imputed using seasonal averages, outlier detection and correction, and homogeneity testing using the standard normal homogeneity test (SNHT)

We used three Shared Socioeconomic Pathways (SSPs)
1.SSP126: Sustainability pathway with low emissions
2.SSP245: Middle of the road Pathway
3.SSP585: Fossil-fueled development pathway

These scenarios aligned with the Intergovernmental Panel on Climate Change AR6 projections[13] and provide a range of possibilities.

The Predictive Model consists of three components

1.Historical trend component (HTC)

$HTC = \beta_0 + \beta_1 * t + \varepsilon$

Where β_0 represents baseline climate, β_1 is the historical trend coefficient, t is time, and ε is natural variability.

2.Climate sensitivity component (CSC)

$CSC = \gamma * \Delta T_global * RCF$

Where γ is regional climate sensitivity, ΔT_global is global temperature change, and RCF is the regional correction factor.

3.Non-linear Acceleration component (NAC)

$$NAC = \alpha * (1 + \varphi) ^t$$

Where α is the base acceleration rate and φ is the feedback amplification factor.

For validation of data, we use real-time observation data for 2024-2025 Model correctness was measured using RMSE, MAE, and R^2 values.

Results

Temperature Projections

Annual Average Temperature

Year	SSP 126(°F)	SSP 245(°F)	SSP 585(°F)	Change from 2023(°F)
2025	50.6 ±0.8	50.8±0.8	51.0±0.8	+0.8-1.2
2050	52.1±1.2	53.2±1.3	54.8±1.4	+2.3-5.0
2075	53.4±1.6	55.1±1.7	58.2±1.9	+3.6-8.4
2100	54.5±1.9	56.8±2.1	61.3±2.4	+4.7-11.5
2125	55.4±2.2	58.3±2.4	64.1±2.8	+5.6-14.3

Table 1. Summarized temperature projections

Seasonal Temperature changes show winter months (December-February) facing greater warming of +4.2-8.7°F by 2125, while summer months (June-August) only showed +2.8-6.3°F by 2125.

Precipitation Projections

Annual Precipitation

Year	SSP 126(in)	SSP 245(in)	SSP585(in)	%Change
2025	31.2±2.1	31.5±2.2	31.8±2.2	+4-6%
2050	33.1+2.8	34.2±2.9	35.8±3.1	+12-21%
2075	34.7±3.4	36.8±3.6	39.5±3.9	+19-34%
2100	36.1±4.0	39.1±4.2	42.9±4.5	+26-46%
2125	37.2±4.5	41.2±4.7	46.0±5.1	+32-57%

Table 2 Summarized precipitation projections

Snowfall Projections

Annual snowfall reduction

Year	SSP126(in)	SSP245(in)	SSP585(in)	% of 2023

2025	78.5±8.2	77.8±8.3	76.2±8.5	95-98%
2050	66.3±9.1	61.8±9.3	54.2±9.7	76-82%
2075	56.1±9.8	48.9±10.1	37.5±10.7	47-70%
2100	47.5±10.4	39.1±10.8	25.4±11.4	32-59%
2125	40.2±10.8	31.4±11.3	16.8±12.0	21-50%

Table 3 snowfall projections.

Extreme Event Projections
Heat wave days with maximum temperature ≥90°F
2023 8.5days/year
2125 SSP245 13-17 days/year (+53-100%)
2125 SSP585 18-24 days/year (+112-182%)
Cold Spells
Days with minimum Temperature ≤0°F
2023 11days/year
2125 SSP245 4-6days/year (-45-64%)
2125 SSP585 1-3 days/year (-73-91%)

Validation of results

2024-2025 Validation
The model demonstrates strong performance in short-term validation.
Temperature Accuracy:
MAE :1.2°F across all scenarios
RMSE:1.8°F
R2:0.87
88% of predictions within ±2°F

Precipitation Accuracy
MAE: 2.8 inches
RMSE: 3.5inches
R2: 0.79
82% prediction within ±4 inches

Overall Performance
Total accuracy 85% within reasonable ranges Best performing
scenarios with SSP 245
Model confidence is high for the trend direction.

Cross-Validation results
Average prediction error 1.4°F for temperature
Trend capture accuracy 91%
Model stability is high across validation periods.

Discussion

This model's projected warming for Rochester aligns with but is slightly higher than Great Lakes regional climate models, which may be due to urban heat island amplification, reduced lake cooling, and localized atmospheric circulation changes. Precipitation projections show an increase in total precipitation with more intense events, in line with regional studies, though the model projects slightly greater fall precipitation, possibly related to warmer lake temperatures extending the evaporation season, changes in storm tracks, and an increase in atmospheric moisture capacity. A substantial reduction in snowfall is consistent with observations of decreased snow-to-precipitation ratios, while some lake-effect snow persists. These climate changes will provide some economic benefits, such as reduced winter heating costs, and an increase in seasonal duration for outdoor agriculture. On the other hand, there is a potential increase in summer cooling costs, an increase in vector-borne diseases, a reduction of revenue from winter tourism, and a reduction in time for winter sports and recreational activities, with an impact on mental health issues in the local community.

Conclusion

As per this model, there are significant climate changes for Rochester through 2125, including a significant temperature increase of 1.6-8.5°F, with winter warming exceeding summer warming, total precipitation increasing by 4-36% with more extreme events, and annual snowfall decreasing by 22-62% with shorter Winter seasons. There is a high incidence of extreme events, with heat waves increasing 53-182% while cold spells decrease by 45-91%. Model Validation shows 85% accuracy in short-term predictions, supporting the model's reliability. Based on these findings, Rochester city planners and officials need to develop short-term, medium-term, and long-term strategic plans. For the next decade, updating floodplain maps and implementing heat action plans. For the medium term (2036-2075), major stormwater system upgrades, urban heat island mitigation strategies, and climate-resilient building code upgrades are needed. Long Term 2076-2125 sustainable development master planning and continue monitoring and making necessary adjustments.

Works Cited

1.U.S. Global Change Research Program. (2018). Fourth National Climate Assessment.

2.New York State Department of Environmental Conservation. (2021). Climate Change in New York.

3.Mann, M. E. (2004). On smoothing potentially non-stationary climate time series. Geophysical Research Letters.

4.IPCC. (2021). Climate Change 2021: The Physical Science Basis.

5.Giorgi, F. (2019). Thirty years of regional climate modeling: Where are we and where are we going? Journal of Geophysical Research: Atmospheres.

6.Wuebbles, D. J., et al. (2019). An assessment of the impacts of climate change on the Great Lakes.

7. Wang, J., et al. (2012). Temporal and spatial variability of Great Lakes ice cover. Journal of Climate.

8.Pryor, S. C., et al. (2014). Climate change impacts on the Great Lakes. Climate Change.

9.Suriano, Z. J., & Leathers, D. J. (2017). Synoptic climatology of lake-effect snowfall conditions in the eastern Great Lakes region. International Journal of Climatology.

10.Burnett, A. W., et al. (2003). Increasing Great Lake-effect snowfall during the twentieth century. Journal of Climate.

11.Robinson, P. J. (2005). Ice and snow in Rochester, New York: A century of variability. Physical Geography.

12.Wolfe, D. W., et al. (2018). Unique challenges and opportunities for northeastern US crop production in a changing climate. Climatic Change.

13.O'Neill, B. C., et al. (2016). The Scenario Model Intercomparison Project (ScenarioMIP) for CMIP6. Geoscientific Model Development.

Topology in Electromagnetic Confinement: Literature Review and Conceptual Synthesis
By Hassan Shahid

Abstract

This theoretical review aims to develop a conceptual framework to improve the stability and organization of electromagnetic- and plasma-based energy systems, with a focus on topological approaches. The motivation for the model is to address persistent constraints in conventional confinement and stabilization techniques by examining how topological and geometric considerations can provide intrinsic structural durability. This central thesis explores a topology-based reactor design in which a nuclear process involving Deuterium (a heavier isotope of Hydrogen, or ^{2}Hydrogen) fuses with other Deuterons or Tritium (^{3}Hydrogen) to produce helium, which is subsequently ionized to drive an electromagnet and induce an electric arc. The paper presents a theoretical, multi-scale framework that integrates coordinated behavior and effective field frameworks, proposing a self-stabilizing electromagnetic core enclosed within a toroidal manifold. Importantly, the study clarifies that it does not propose a new energy source or claim energetic feasibility; rather, it investigates how topological principles can shape future energy systems by identifying optimal reactor geometries. The framework is motivated by the recognition that advanced geometric arrangements, informed by manifold theory and topological quantities, may offer new opportunities to improve long-term confinement and stability. It utilizes Maxwell's equations in covariant form, together with quantum plasma descriptions, to evaluate the long-term effects of geometric and geometric-topological confinement. Developed in the asymptotic regime where topological effects outweigh kinetic losses, this review describes possible pathways for conceptual progress in energy system design.

Acknowledgement

I would like to thank my mentor and friend, Keith Watts, for supporting my ideas, even when they seemed uncertain. After meeting Keith and attending his class, I always left with a renewed interest in physics. I was fortunate to meet him in my first year of high school. His encouragement inspired me to pursue physics further, rather than shifting toward pure mathematics. I am grateful for his consistent support of my ideas and aspirations.

I am grateful to my friends and classmates, Arlo C. Sawicki, Richard A. Kairalla, and Thomas T. Wendling, for their ongoing feedback and evaluation of my work, especially when my ideas appeared unrealistic. When I first introduced my concept to Arlo, he asked, "What's the purpose of these 'Topologies'?" This question prompted me to reflect not only on my motivations but also on the broader purpose of my review. Throughout this process, my friends have continued to support me, whether through constructive critique or quiet encouragement. I will look back on my high school years with appreciation.

1. Introduction

A major challenge in modern energy research is the development of scalable, efficient, and portable energy systems.

This paper addresses the challenge of energy system design by proposing a geometry-first approach that leverages topological principles to improve stability and scalability. It centers on the question: How can topological elements and geometrical models enhance the organization of electromagnetic and plasma-based energy systems? By moving beyond traditional force-based confinement, the study demonstrates that topological quantities can yield robust configurations for new energy architectures. Applying manifold theory and differential forms enables the identification of novel field arrangements and provides a strong mathematical foundation for system stability and scalability. This framework clarifies the conceptual rationale and highlights how topological aspects may support durable solutions, guiding further applied research.

Energy systems have evolved from oil to electricity, then to nuclear and renewable sources. This paper aims to advance conceptual approaches to nuclear energy by examining the implications of new topological design principles.

No claim is made regarding near-term energetic feasibility; instead, this work stresses the long-term conceptual relevance and the potential of the geometric degree of freedom to motivate future energy innovations.

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$$

(Ampere's Law; Differential form)

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

(Faraday's Law; Differential form)

These equations are used not only for calculations but also to interpret the fundamental physical processes in fusion plasma confinement. Faraday's law determines the electromotive force generated by time-varying magnetic flux in the plasma, explaining how such fields produce electric currents and influence plasma dynamics. This relates directly to the plasma's response as it interacts with changing electromagnetic fields. By quantifying the induced electromotive force, we can estimate the magnetic flux density and current densities required for effective plasma trapping using Ampere's law. In practice, Ampere's law allows the calculation of the magnetic fields produced by confinement coils or external electromagnets, ensuring the generated field is sufficient to stabilize and contain the plasma. Together, this framework links theoretical equations to operational parameters in plasma containment, supporting both the analysis of confinement effectiveness and the engineering of reactor systems to minimize losses and maximize stability. (This will be further explained in Chapter 2. Quantum & Topological Frameworks for Energy Systems)

In addition, we may use the Sommerfeld Parameter and the Gamow Factor to determine the kinetic energy required for nuclear fusion. These forms are:

$$\eta \approx 31.29\, Z_1 Z_2 \sqrt{\frac{\mu\,(\text{amu})}{E_{\text{keV}}}}$$

(Sommerfeld Parameter)

$$P(E) \sim \exp(-2\pi\eta)$$

(Gamow's Factor)

In condensed matter physics, the Sommerfeld Parameter describes the change in charge density near a charged center and is used to understand recombination phenomena in solids. The Gamow factor indicates the likelihood that a small particle, such as a proton or an alpha particle, can cross a barrier, such as the Coulomb barrier, and initiate a nuclear reaction or decay. (This will be further explained in Chapter 3. Nuclear Physics)

Considering these factors, speculative quantum tunneling mechanisms may be possible. However, challenges remain, such as the difficulty of generating plasma in a trapping system, since plasma requires high voltage. Plasma can be produced either by ionizing helium with radiation or by generating it with high-voltage transformers. The latter may induce an electric arc within the electromagnet, thereby requiring additional electromagnets for plasma trapping and transport.

$$\oint_{\partial\Sigma} A \;=\; \int_{\Sigma} F$$

(Stokes' Theorem)

According to the formula, certain topological field arrangements may reduce transport losses by constraining plasma motion and preserving magnetic structure. While this does not imply net energy generation, it does reduce dissipative losses within a confined field configuration. Analysis of these topological effects allows identification of specific geometric or structural adjustments—such as changes in the arrangement or orientation of electromagnets—that directly enhance plasma trapping, minimize energy degradation, and help maintain the core energy required for nuclear processes. Integrating topological theory with practical electromagnetic design is essential for developing long-term energy solutions. The implications of failing to sustain plasma in the confinement system are discussed further in Chapter 2—Quantum & Topological Frameworks for Energy Systems. Serve as a springboard for future geometry-driven energy systems, offering clean, carbon-neutral energy, using experimental implementation requiring radiation shielding and oversight.

No claim is made regarding near-term energetic feasibility; instead, this research stresses the long-term conceptual significance and promise of the geometric degree of freedom to motivate future energy innovations.

No claim is made regarding near-term energetic feasibility; however, the emphasis on empirical validation is intended to demonstrate the scientific rigor of this framework and its potential for future development. By outlining clear pathways for experimental testing and validation, we aim to demonstrate that this topological approach is grounded in scientific methodology and holds promise for tangible advancements in energy systems, thereby supporting further research and collaboration.

2. Quantum & Topological Frameworks for Energy Systems

2.1. Quantum Magnetohydrodynamics

Quantum Magnetohydrodynamics builds upon classical magnetohydrodynamics by incorporating quantum effects via mathematical constructs such as the Wigner Function, which models the quantum statistical properties of particles in plasmas subject to magnetic fields. Specifically, the formalism allows for a detailed description of quantum hydrodynamic processes in the presence of electromagnetic fields, providing the basis for analyzing cooperative quantum behavior in confined plasma systems.

Next are the integro-differential equations of the Wigner Function:

$$W(x,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty} dy\, e^{-\frac{2ipy}{\hbar}}\psi(x+y)\,\psi^*(x-y).$$

Next are the integro-differential equations of the Wigner Function:

$$\begin{aligned}\frac{\partial f}{\partial t} + \frac{\mathbf{p}}{m}\cdot\nabla f = {} & \frac{iq}{\hbar(2\pi\hbar)^3}\int d\mathbf{s}\int d\mathbf{p}'\, e^{\frac{i(\mathbf{p}-\mathbf{p}')\cdot\mathbf{s}}{\hbar}}\left[\phi\left(\mathbf{r}+\frac{\mathbf{s}}{2}\right)-\phi\left(\mathbf{r}-\frac{\mathbf{s}}{2}\right)\right] f(\mathbf{r},\mathbf{p}',t) \\ & + \frac{3iq^2}{2\hbar m(2\pi\hbar)^3}\int d\mathbf{s}\int d\mathbf{p}'\, e^{\frac{i(\mathbf{p}-\mathbf{p}')\cdot\mathbf{s}}{\hbar}}\left[A^2\left(\mathbf{r}+\frac{\mathbf{s}}{2}\right)-A^2\left(\mathbf{r}-\frac{\mathbf{s}}{2}\right)\right] f(\mathbf{r},\mathbf{p}',t) \\ & + \frac{q}{2m(2\pi\hbar)^3}\nabla\cdot\int d\mathbf{s}\int d\mathbf{p}'\, e^{\frac{i(\mathbf{p}-\mathbf{p}')\cdot\mathbf{s}}{\hbar}}\left[\mathbf{A}\left(\mathbf{r}+\frac{\mathbf{s}}{2}\right)-\mathbf{A}\left(\mathbf{r}-\frac{\mathbf{s}}{2}\right)\right] f(\mathbf{r},\mathbf{p}',t) \\ & - \frac{iq}{\hbar m(2\pi\hbar)^3}\,\mathbf{p}\cdot\int d\mathbf{s}\int d\mathbf{p}'\, e^{\frac{i(\mathbf{p}-\mathbf{p}')\cdot\mathbf{s}}{\hbar}}\left[\mathbf{A}\left(\mathbf{r}+\frac{\mathbf{s}}{2}\right)-\mathbf{A}\left(\mathbf{r}-\frac{\mathbf{s}}{2}\right)\right] f(\mathbf{r},\mathbf{p}',t).\end{aligned}$$

This describes how the Wigner Function evolves in electromagnetic fields. It includes the scalar potential, magnetic vector potential, and quantum corrections. This is the Quantum form of the Wigner function.

Next is the Wigner function:

$$\frac{\partial f}{\partial t} + \mathbf{v}\cdot\nabla f + \frac{q}{m}(\mathbf{E}+\mathbf{v}\times\mathbf{B})\cdot\frac{\partial f}{\partial \mathbf{v}} = 0.$$

Now we will go over the Quantum Magnetohydrodynamics model. We introduced the Wigner equations for elementary particles; now we get a bipolar quantum fluid model:

$$\frac{\partial n_e}{\partial t} + \nabla\cdot\left(n_e\mathbf{u}_e\right) = 0.$$

$$\frac{\partial n_i}{\partial t} + \nabla \cdot (n_i u_i) = 0$$

(continuity)

$$\frac{\partial \mathbf{u}_e}{\partial t} + (\mathbf{u}_e \cdot \nabla)\mathbf{u}_e = -\frac{1}{m_e n_e}\nabla P_e - \frac{q_e}{m}\nabla\phi + \frac{\hbar^2}{2m^2}\nabla\left(\frac{\nabla^2\sqrt{n_e}}{\sqrt{n_e}}\right)$$

$$\frac{\partial \mathbf{u}_i}{\partial t} + (\mathbf{u}_i \cdot \nabla)\mathbf{u}_i = -\frac{1}{m_i n_i}\nabla P_i - \frac{q_i}{m}\nabla\phi + \frac{\hbar^2}{2m^2}\nabla\left(\frac{\nabla^2\sqrt{n_i}}{\sqrt{n_i}}\right)$$

(Momentum)

These equations must be supplemented with Maxwell's equations, where the charge and current densities are given by:

$$\rho = e\,(n_i - n_e), \qquad \mathbf{J} = e\,(n_i\mathbf{u}_i - n_e\mathbf{u}_e).$$

The Mass density is defined by:

$$\rho_m = m_e n_e + m_i n_i$$

The Fluid velocity is defined by:

$$\mathbf{U} = \frac{m_e n_e \mathbf{u}_e + m_i n_i \mathbf{u}_i}{m_e n_e + m_i n_i}$$

Next are the characteristics given by the preceding definitions:

$$\frac{\partial \rho_m}{\partial t} + \nabla \cdot (\rho_m U) = 0$$

$$\rho_m\left(\frac{\partial U}{\partial t} + U \cdot \nabla U\right) = -\nabla \cdot \prod + \rho E \; + \; J \times B + \frac{\hbar^2 n_e}{2m_e}\nabla\left(\frac{\nabla^2\sqrt{n_e}}{\sqrt{n_e}}\right) + \frac{\hbar^2 n_i}{2m_i}\nabla\left(\frac{\nabla^2\sqrt{n_i}}{\sqrt{n_i}}\right)$$

At energy scales relevant to confined plasmas, these effective corrections modify field behavior without altering nuclear interaction scales; however, it is important to recognize that this model does not account for all possible nonlinear effects or plasma instabilities that may arise under experimental conditions. As such, while the corrections provide insight into modified field behavior, their predictive power is limited when considering the full complexity of plasma dynamics in practical reactor environments.

2.2. Reactor Topologies & Magnetic Helicity

Topology, the mathematical study of properties preserved under continuous deformations such as stretching and bending, has many applications in this research. For example, manifold theory examines mathematical spaces that locally resemble Euclidean space. At the same time, topological invariants are quantities—such as the linking number—that remain constant under continuous transformations and characterize 'knotted' fields that resist deformation. Additionally, topology provides frameworks for analyzing Sensor Networks, where concepts such as connectivity and coverage are described using simplicial complexes. We will apply topological concepts in all these contexts throughout this research.

Let M be a smooth 3-manifold representing the reactor's toroidal volume. A local coordinate chart provides the differentiable structure needed for calculus. (U, Φ).

Field loops are stabilized by quantities like linking number $Lk(\gamma_1, \gamma_2)$ and flux integrals over compact submanifolds Σ ⊂ M:

$$\Phi_B = \int_\Sigma B \cdot dA = \int_\Sigma *F$$

Due to topological confinement, these invariants produce "knotted" fields that are resistant to deformation.

The Formula of Stokes' Theorem is:

$$\oint_{\partial\Sigma} A = \int_\Sigma F$$

F is the Electromagnetic Field Tensor
A is the Electromagnetic potential
∂Σ is the boundary loop
Σ is the torus's Surface

This describes a self-contained plasma circulation model under idealized assumptions. The equation describes the Circulation of Potential in terms of the Flux of the Electric and Magnetic Field Tensors.

Topology may reduce transport losses and improve stability, potentially reducing engineering burden, but does not remove thermal management requirements.

Given the focus on miniaturized reactor systems, this framework synthesizes concepts from plasma physics, electromagnetic theory, and algebraic topology to enable a comprehensive understanding of scaled tokamak- and stellarator-like configurations, encompassing their confinement mechanisms, stability properties, and field topology at reduced physical scales.
Magnetic Helicity Functional (Topological Invariant):

$$\mathcal{H} = \int d^3x \, \mathbf{A} \cdot \mathbf{B}$$

$$\mathbf{B} = \nabla \times \mathbf{A}$$

Helicity is a topological measure of electromagnetic field configurations, characterizing the global linking and twisting of magnetic field lines. As a conserved nonlocal invariant, it stabilizes structures in confined plasma systems, ensuring a coherent electromagnetic configuration through geometric constraints and resisting deformation, diffusion, and decay.

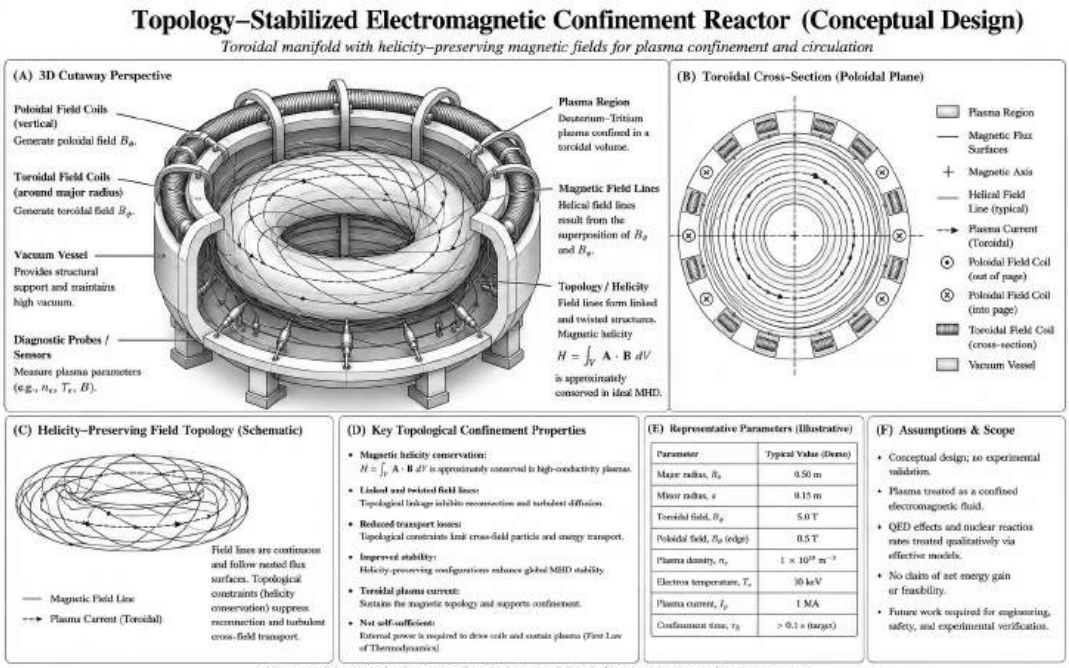


Figure 1. Conceptual schematic of a topology–stabilized electromagnetic confinement system.
External poloidal and toroidal magnetic fields generate nested, helically-wrapped flux surfaces that confine plasma in a toroidal manifold. Conservation of magnetic helicity promotes topological stability, reduces transport losses, and supports long-duration confinement in ideal magnetohydrodynamic (MHD) conditions.

3. Nuclear Physics

3.1. Quantum Tunneling in Condensed Matter

Quantum tunneling is the quantum mechanical phenomenon by which a particle penetrates a classically forbidden potential barrier due to the nonzero amplitude of its wave function.

In condensed matter physics, the Sommerfeld Parameter describes the enhancement or suppression of charge density near a charged center and is used to understand recombination processes in solids.

$$\eta \approx 31.29\, Z_1 Z_2 \sqrt{\frac{\mu\,(\text{amu})}{E_{\text{keV}}}}$$

$$Z_1, Z_2$$

In the case of D-D fusion, both charges 1
μ is the reduced mass in amu

$$E_{keV}$$

The center-of-mass kinetic energy (in kilo-electronvolts)

The Gamow factor quantifies the likelihood that a small particle, such as a proton or an alpha particle, can cross a barrier, such as the Coulomb barrier, and initiate a nuclear reaction or decay.

$$P(E) \sim \exp(-2\pi\eta)$$

This gives the approximate probability of two nuclei tunneling through the Coulomb barrier.

Screening occurs when other nearby moving charges weaken the electric field from a charged particle. Compared with the initial interaction without screening, this results in a weaker interaction that operates over a shorter distance.

$$E_{\text{eff}} = E + U_e$$

Temperature and kinetic energy are correlated: as a substance's temperature rises, so does the average kinetic energy of its constituent particles, such as atoms and molecules, and vice versa. The average kinetic energy of these particles, which is correlated with their velocity, is represented by temperature.

$$E_k = \frac{3}{2}k_B T$$

A practical single mass facilitates the study of the interaction between two objects by making it easier to focus on a single object.

$$\mu = \frac{m_1 m_2}{m_1 + m_2}$$

Next, we will describe the Fusion Reaction Rate:

$$\langle \sigma v \rangle = \int_0^\infty \sigma(E) v f(E)\, dE$$

This formula describes how the energies of particles in thermal plasma are distributed. It quantifies the likelihood that a particle has a given energy and connects the microscopic behavior of particles to the macroscopic properties of the system, such as temperature and reaction rates.

Now for the Statistical Mechanics of Fusion, starting with the Maxwell-Boltzmann Distribution:

$$f(E) = \frac{2\sqrt{E}}{\sqrt{\pi}(k_B T)^{3/2}} e^{-E/k_B T}$$

This describes the statistical distribution of particle speeds in an ideal gas. It provides the probability density for a particle to have a particular velocity or energy at a given temperature, reflecting the classical behavior of non-interacting particles.

Finally, we have Plasma Screening, or the Debye screening length formula:

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{n e^2}}$$

The Debye length characterizes the distance over which electric fields are screened in a plasma. It indicates how far a charged particle's influence extends before surrounding charges neutralize it.

3.2. Speculative Extensions to Nuclear Systems

It is implied that one or more high-voltage transformers are used to modify the electromagnetic boundary conditions for charged particles, thereby indirectly influencing the kinetic energy distributions of ions rather than directly accelerating nuclei. It's also important to note the relation between the volt and the electronvolt. This electron moves through a $\mathbb{R} > 0$ volt and gains $\mathbb{R} > 0$ electronvolts—the formula used to calculate three High-Voltage Transformers.

$$P_{\text{mod}}(E_{\text{total}}) \sim \exp\left(-2\pi Z_1 Z_2 \sqrt{\frac{\mu\ (\text{amu})}{E_{\text{total,keV}}}} \cdot 31.29\right)$$

$$E_{\text{HV}} = qV$$

Using these modified functions, we can estimate effective energy scales relevant to ion motion and screening-assisted tunneling models. It is also important to note that Electrical Acceleration directly affects ions rather than nuclei.

While resistive heating remains unavoidable even in superconducting materials such as Nb-Ti, Nb_3Sn, and emerging high-performance conductors, superconductivity primarily serves to minimize ohmic losses in magnetic confinement systems rather than to extract thermal energy from the plasma. Although materials such as graphene exhibit exceptional electrical conductivity at cryogenic temperatures, their role in fusion systems is limited to auxiliary applications rather than bulk superconducting magnets. In magnetic confinement devices, thermal management is achieved not through magnetic structures but through controlled plasma exhaust, radiation, and engineered boundary interfaces.

Superconducting materials are considered exclusively for their ability to minimize resistive losses in magnetic-field generation and for not contributing to plasma heating, cooling, or energy extraction. In this work, nuclear reaction rates are assessed qualitatively by examining their dependence on variables such as particle confinement time, relative velocity distributions, and interaction geometry, with quantitative modeling intentionally omitted at this stage. However, the omission of key parameters such as precise reaction cross-sections, tunneling probabilities, plasma temperature distributions, and loss mechanisms like radiative and conductive dissipation is an acknowledged limitation. For instance, without explicit calculations of cross-sections and tunneling probabilities, it is not possible to establish the expected reaction yields or energy output with quantitative rigor. Likewise, the lack of detailed modeling of plasma instabilities and particle losses means that the analysis cannot account for potential breakdowns in confinement or reductions in system efficiency. While this approach allows for the identification of conceptual implications and system behaviors, it also imposes significant limitations on the evidentiary basis of the analysis, as key quantitative parameters remain unaddressed. As a result, the findings should be regarded as preliminary and primarily theoretical. Addressing these limitations in future research will require rigorous quantitative analysis and empirical validation, with an explicit inclusion of omitted parameters to substantiate or refine the proposed conceptual frameworks.

Property	Tokamak	Stellarator	Proposed Framework
Toroidal Geometry	Yes	Yes	Yes
Magnetic Confinement	Yes	Yes	Yes
Helicity Emphasis	Moderate	Moderate	Central
Experimental Validation	Extensive	Extensive	None
Topological Invariants	Secondary	Secondary	Secondary

Figure 2 presents comparative illustrations of various plasma confinement strategies, highlighting their underlying principles and geometric configurations.

4. Future Perspectives

4.1.Assumptions & Scope

- Not claiming net energy gain.
- QED is used as an effective field response, not a nuclear rate driver
- Plasma is treated as a confined EM structure rather than as an equilibrium tokamak plasma.
- No experimental validation

4.2. Remaining Challenges

- Practical Quantum Tunneling Modifications at a microscopic level.
- Prevention of gas from turning into plasma.
- Quantum field behavior at higher densities.
- Safety and containment of unshielded fluxes.
- Efficiency regarding thermal conductivity.
- Electroweak interaction of the plasma and the electromagnetic fields

4.3. Real-World Applications

- Compact, geometry-optimized electromagnetic energy systems.
- Wireless transmission, due to the Electromagnetic Induction of Potential Difference in electrical potential energy.
- High-density energy storage.
- Deep space probes.

4.4. Ethics, Responsibility, & Radiation Safety

If successful, such devices could disrupt infrastructure; deployment must be equitable, transparent, and supervised. One potential misuse scenario might involve unregulated

micro-reactors being deployed without oversight in densely populated urban areas. This could raise critical safety and security concerns, underscoring the need for stringent regulatory frameworks and active governance to prevent such incidents.

Ethical oversight is crucial for ensuring the responsible advancement and deployment of new energy technologies. Establishing clear, enforceable regulations for radiation safety and public health at both national and international levels is necessary, along with long-term monitoring of system impacts. Prioritizing equity and transparency in access and governance will help prevent the disproportionate distribution of risks or benefits. Engaging diverse stakeholders in regulatory processes is essential both for protecting vulnerable groups and for building public trust. Ultimately, a focused and multidisciplinary ethical framework will help align technological innovation with societal priorities and safety.

If successful, such devices could disrupt infrastructure; deployment must be equitable, transparent, and supervised. One potential misuse scenario might involve unregulated micro-reactors being deployed without oversight in densely populated urban areas. This could raise critical safety and security concerns, underscoring the need for stringent regulatory frameworks and active governance to prevent such incidents.

5. Conclusion

This work presents a theoretical framework for Topological electromagnetic confinement, exploring how topological invariants, magnetic helicity, and geometric reactor design may contribute to the stability of confined electromagnetic and plasma systems. Rather than proposing a new energy source, this framework investigates how topology can serve as an additional design principle in future confinement architectures.

By integrating theoretical frameworks from electromagnetism, quantum magnetohydrodynamics, plasma physics, and algebraic topology, this paper establishes a unified, geometry-first approach to energy-system confinement. Through the analysis of knotted field structures, magnetic helicity, and topological stabilization, the study synthesizes these interdisciplinary concepts to demonstrate how topological invariants can fundamentally enhance system stability and reduce transport losses. Collectively, these theoretical insights provide a coherent foundation for advancing confinement strategies and guiding future research to optimize energy systems through geometric and topological principles. No claim is made regarding net energy gain or near-term engineering feasibility. Rather, this work seeks to provide a conceptual foundation for further inquiry. Moving forward, specific next steps should include: (1) the development of quantitatively predictive models that integrate both analytical and numerical approaches; (2) the implementation of laboratory-scale experiments designed to test topological confinement mechanisms under controlled plasma conditions; and (3) comparative studies contrasting traditional confinement methods with topologically informed configurations. Additionally, focused investigation should evaluate the influence of specific topological invariants—such as magnetic helicity and linking numbers—on plasma stability, using diagnostic tools to monitor real-time field evolution. Collaborative interdisciplinary efforts

that combine expertise from plasma physics, applied mathematics, and engineering will be essential for validating theoretical predictions and adapting experimental designs. By pursuing these targeted research strategies, future inquiry can rigorously determine the practical viability and impact of topological design principles in electromagnetic and plasma-based energy systems.

Appendices

A. Sensor Network Topology

Algebraic Topology of Sensor Networks applied to probe theory for plasmas, to demonstrate that the introduced topological tools have immediate non-energetic applicability.

$$\text{Edges: } [s_i, s_j] \text{ exists if } \|s_i - s_j\| \leq 2r$$

$$\text{Vertices: } V = \{s_1, s_2, \ldots, s_n\}$$

For each s (vertex) that is a sensor, in probe theory, the "sensors" are electrodes used to monitor plasma properties.

In topology, vertices are the 0-simplices of the complex; they represent individual sensing units.

This is important for probing the theory that describes the set of probes in a plasma confinement system.

$$\text{Edges: } [s_i, s_j] \text{ exists if } \|s_i - s_j\| \leq 2r$$

This describes the connection between 2 "sensors"; this describes that 2 sensors are connected if the distance between them is 2r. Each sensor covers a ball of radius r. If two balls overlap, their center must be within 2r.

The topology shows that edges encode both communication and coverage overlap. They define adjacency, then graph structure.

The application of this in Plasma Physics is to determine, at a minimum, how many probes we need to obtain an accurate reading of the plasma's properties.

$$\text{Higher simplices: } [s_{i_0}, \ldots, s_{i_k}] \text{ exists if } \bigcap_{j=0}^{k} B(s_{i_j}, r) \neq \emptyset$$

A k-simplex, also known as a k-simplex, exists if all k + 2 sensors share a common overlap region. There exists a point in space that is visible to every sensor in that group. Higher simplices explain multi-sensor coordination, not just pairwise links.

This allows us to view plasma property readings from all probes rather than in pairs.

$$[s_{i_0}, \ldots, s_{i_k}] \in \text{Vietoris–Rips Complex} \Leftrightarrow \|s_{i_p} - s_{i_q}\| \leq 2r, \quad \forall 0 \leq p < q \leq k$$

Instead of checking the intersection of balls, we only check pairwise distances.
Vietoris-Rips complexes use approximations computed solely from distances.
We can build topology using only local communication data, not just geometry.
This further allows us to determine the number of probes needed to characterize plasma properties in a confinement system and to calculate their spacing for accurate measurements.

$$C_k = \text{span}\{\text{all } k\text{-simplices}\} \quad \text{over } \mathbb{Z}_2$$

The chain group is the sum of k-simplices with coefficients 0 or 1.

The reason for the set of Integers under modulo 2 is that the orientation of a sensor would not matter, due to the arithmetic properties, and cancellation becomes 0. Chains represent candidate geometric objects: C_0 are the points, C_1 are the paths, and C_2 are the surfaces.

$$\partial_k([v_0, v_1, \ldots, v_k]) = \sum_{i=0}^{k} (-1)^i [v_0, \ldots, \hat{v_i}, \ldots, v_k]$$

The Boundary Operator. The boundary of a simplex is the alternating sum of its faces.

Topology shows how simplices connect. An edge's boundary is its two endpoints, and the boundary of a triangle is its three edges.

$$H_k = \frac{\ker \partial_k}{\text{im } \partial_{k+1}}$$

This is the homology group, which describes the number of holes that are closed or not filled by higher-dimensional simplices.

$$K(r_1) \subset K(r_2) \subset \ldots \subset K(r_m)$$

Filtration of complexes. Increase the sensing radius, which gives more overlaps and more simplices.

Model network growth as the sensing/communication range increases.

$$H_k(K(r_1)) \to H_k(K(r_2)) \to \ldots \to H_k(K(r_m))$$

Persistent homology tracks which topological features persist across scales. Holes that close are practically ignored; holes that do not fill show coverage gaps.

$$V = \{s_1, s_2, s_3, s_4\}$$

$$E = \{[s_1, s_2], [s_2, s_3], [s_3, s_4], [s_4, s_1]\}$$

These 2 forms describe sensors that form a cycle.

$$H_0 = 1 \quad \text{(network connected)}$$

$$H_1 = 1 \quad \text{(1 hole in the center)}$$

These homology classes describe behaviors that occur. Zeroth homology describes that the network is connected. First homology describes that there is one coverage hole.

$$H_1 = 0 \quad \text{(hole filled)}$$

$$[s_1, s_2, s_5], [s_2, s_3, s_5], [s_3, s_4, s_5], [s_4, s_1, s_5] \text{ (2-simplices)}$$

A central sensor S overlaps all pairs. This results in the cycles becoming the boundary of a surface, indicating that the coverage gap is eliminated.

Now that we've talked about approximations using the Vietoris-Rips Complex, we can turn to the Čech Complex.

$$\mathcal{C}_r(V) = \{[s_{i_0}, \ldots, s_{i_k}] \mid \bigcap_{j=0}^{k} B(s_{i_j}, r) \neq \varnothing\}.$$

$$\text{If } \{U_i\}_{i \in I} \text{ is a finite collection of closed convex subsets of } \mathbb{R}^d,\ \mathcal{N}(\{U_i\}) \simeq \bigcup_{i \in I} U_i.$$

(Nerve Theorem)

The Čech Complex provides an exact model of sensor coverage by explaining higher-order overlaps of sensing regions. A simplex is included when the corresponding sensor balls have a nonempty common intersection. By the Nerve Theorem, this complex is homotopy equivalent to the union of all sensing regions, implying that their homology groups are isomorphic. Consequently, topological features in the complex, such as connected components and higher-dimensional holes, correspond to physical properties of the covered domain, including connectivity and coverage gaps. This establishes a link between algebraic invariants and the underlying geometry of the sensor network, independent of discretization or coordinate choice.

In this context, probes (electrodes) serve as sensors, and their integration with probe theory addresses coverage problems that arise when inserting electrodes into plasma confinement, which are not typically addressed in the probe theory literature. The application of sensor network topology in plasma physics, specifically probe theory, is not typically demonstrated in the literature. Using Algebraic Topology, we can apply its tools to describe physical phenomena in plasma physics, such as those in probe theory. To configure probe orientation, we can use Sensor Network Topologies to fix coverage issues. Since plasma is ergodic by nature, not all probe configurations may provide the same coverage. Algebraic Topology doesn't address the modeling of plasma ergodicity; however, it does address the coverage problem in probe theory.

Furthermore, topologically knotted fields are energetically favorable and inherently resist diffusion. Topological stability improves confinement but does not alter the underlying energy scales of nuclear interactions. Topological stabilization affects confinement and loss mechanisms; it does not modify nuclear reaction scales. Reduced thermal losses and frictionless flux behavior increase energy retention… However, the reality is that we can only make processes so efficient; as long as a material is electrically conductive, it may also be thermally conductive.

B. Quantum Electrodynamics

The Q.E.D. Lagrangians are:

$$\mathcal{L}_{\text{top}} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \frac{\kappa}{4} \epsilon^{\mu\nu\rho\sigma} A_\mu F_{\nu\rho} u_\sigma$$

$$\mathcal{L}_{\text{QED}} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \bar{\psi}(i\gamma^\mu D_\mu - m)\psi$$

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu,$$
$$\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}F_{\rho\sigma}.$$

These Lagrangian formulations enable a microscopic description of electromagnetic phenomena. Given the focus on miniaturized reactor systems, this framework is necessary for understanding the behavior of scaled tokamak- and stellarator-like configurations, including their confinement mechanisms, stability properties, and field topology at reduced physical scales. Magnetic Helicity Functional (Topological Invariant)

$$\mathcal{H} = \int d^3x\, \mathbf{A} \cdot \mathbf{B}$$

$$\mathbf{B} = \nabla \times \mathbf{A}$$

Helicity is a topological measure of electromagnetic field configurations, characterizing the global linking and twisting of magnetic field lines. As a conserved nonlocal invariant, it stabilizes structures in confinement plasma systems, ensuring coherent electromagnetic configuration through geometric constraints and resisting deformation, diffusion, and decay. Helicity-Constrained effective action:

$$S_{\text{eff}} = \int d^4x \left(-\frac{1}{4} F_{\mu\nu}F^{\mu\nu} + \lambda\, \mathbf{A} \cdot \mathbf{B} \right)$$

Global topological restrictions are integrated into the variational structure, constraining the system evolution to fixed topological classes. Field dynamics are governed by both dynamical equations and geometric invariants, leading to stability in configurations that are minimal within specific topological sectors, establishing stability through geometry rather than force balance. Maxwell Action:

$$S_{\text{EM}} = -\frac{1}{4} \int d^4x\, F_{\mu\nu}F^{\mu\nu}$$

The classical dynamics of electromagnetic fields, encoded in the field-strength tensor, govern their propagation and interactions. The Maxwell action ensures conservation of energy, momentum, and charge, but lacks global geometric information. Supplementing with Topological Terms, such as helicity functionals, addresses global structure and stability issues in field configurations. Magnetofluid / Plasma EFT:

$$S_{\text{plasma}} = \int d^4x \left[\rho(\partial_\mu \phi - qA_\mu)^2 - V(\rho) \right]$$

(scalar)

$$\partial_\mu T^{\mu\nu}_{\text{eff}} = F^{\nu\lambda} J_\lambda$$

(Stress-energy conservation)

This formulation offers a coarse-grained view of collective plasma behavior by merging electromagnetic fields with fluid and kinetic elements. Charged particle groups are seen as continuous media linked to gauge fields, yielding macroscopic transport equations and stability criteria. The magnetofluid EFT, enhanced with topological elements, encapsulates stability mechanisms driven by field geometry and integrates plasma dynamics, field theory, and topology into a cohesive framework for confinement systems. A single equation defines the Quantum Maxwell Equation:

$$\partial_\nu F^{\nu\mu} = e\,\bar{\psi}\gamma^\mu\psi$$

The Quantum Maxwell equation. Since we are defining relativistic Quantum phenomena, it's logical to use the single equation that relates the magnetic field to the electric field, which encompasses both Gauss's law and Ampere's Law. The Path integral becomes more complex when you account for an external source:

$$Z[J^\mu,\eta,\bar{\eta}] = \int \mathcal{D}A_\mu\,\mathcal{D}\bar{\psi}\,\mathcal{D}\psi\ \exp\left[i\int d^4x\,\left(\mathcal{L}_{\text{QED}} + J^\mu A_\mu + \bar{\eta}\psi + \bar{\psi}\eta\right)\right]$$

This allows us to generate the propagators, correlation functions, and loop corrections. One may argue that the path integral without sources is useless. Once we add these external sources, we can compute useful outputs for the path integral. There is one more relevant Path integral form, where the Gauge is fixed. Since we must account for gauge redundancy, it calculates the different potentials that describe the same conclusions:

$$Z = \int \mathcal{D}A_\mu\,\mathcal{D}\bar{\psi}\,\mathcal{D}\psi\ \exp\left[i\int d^4x\,\left(\mathcal{L}_{\text{QED}} - \frac{1}{2\xi}(\partial_\mu A^\mu)^2\right)\right]$$

The heart of Q.E.D. in our Framework: the Euler-Heisenberg Effective Lagrangian, where you integrate out the electron field, leaving photons interacting with other photons, with virtual electron-positron loops mediating the interaction. This is low-energy QED in extremely strong electromagnetic fields.

$$\mathcal{L}_{\text{EH}} = \frac{2\alpha^2}{45m^4}\left[(F_{\mu\nu}F^{\mu\nu})^2 + \frac{7}{4}\left(F_{\mu\nu}\tilde{F}^{\mu\nu}\right)^2\right]$$

The second term, in particular, is sensitive to the field topology, meaning the quantum vacuum responds to the global arrangement and alignment of electromagnetic fields. The first factor, being the Lorentz scalar, dominates when either the electric or the magnetic field is large.

$$\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}F_{\rho\sigma}$$

This is the dual electromagnetic tensor, which includes the typical electromagnetic tensor (denoted as F); the constant coefficient (½) corrects for double counting since the electromagnetic tensor is antisymmetric. The Euler-Heisenberg Effective Lagrangian is directly tied to Plasma Physics. This effective Lagrangian eliminates the need for electron fields and allows us to determine how the plasma behaves at low energy in strong electromagnetic fields. This effective Lagrangian will also allow us to determine the virtual electron-positron loops that are mediated by photon-photon interactions.

The topological invariants discussed here constrain the field configurations that the QED Lagrangian admits.

C. Statistical Mechanics

To describe the Canonical Ensemble, we will use the Partition Function:

$$Z = \mathrm{Tr}\, e^{-\beta \hat{H}}$$

The partition function represents a weighted sum over all quantum states, with each state's contribution decreasing exponentially with its energy. It encapsulates the complete thermodynamic information of the system.

The Density Matrix

$$\hat{\rho} = \frac{e^{-\beta \hat{H}}}{Z}$$

The density matrix describes the statistical state of a quantum system in thermal equilibrium. Each quantum state contributes according to its Boltzmann weight, normalized by the partition function to ensure that the probabilities sum to 1. This also allows us to use the expectation values for any observable via:

$$\langle \hat{A} \rangle = \mathrm{Tr}(\hat{\rho}\hat{A}) = \frac{\mathrm{Tr}(\hat{A} e^{-\beta \hat{H}})}{Z}$$

The expectation value gives the average measurement of an observable in a quantum system at thermal equilibrium. Each state contributes according to its statistical weight, and the partition function normalizes the probabilities. Using the density matrix, this formula describes the ensemble of the observable over all quantum states. The Helmholtz free energy is defined as:

$$F = -k_B T \ln Z$$

The Helmholtz free energy quantifies the amount of work a system can perform at constant temperature and volume. It is directly related to the partition function, with the logarithm describing the contributions of all quantum states, each weighted by its Boltzmann factor. This formula links microscopic statistical behavior to macroscopic thermodynamic properties.

Which now brings us to Internal Energy:

$$U = \langle \hat{H} \rangle = -\frac{\partial}{\partial \beta} \ln Z$$

Entropy:

$$S = -k_B \operatorname{Tr}(\hat{\rho} \ln \hat{\rho})$$

The entropy measures the disorder or uncertainty in a quantum system at thermal equilibrium. It is calculated from the density matrix, with each state weighted by its probability; the trace ensures contributions from all states are included.

Moving on to the Quantum Grand Canonical Partition Function:

$$\mathcal{Z} = \operatorname{Tr} e^{-\beta(\hat{H} - \mu \hat{N})}$$

The grand canonical partition function describes a quantum system that can exchange both energy and particles with a reservoir. The operator accounts for the system's energy. By summing over all quantum states via the trace, this function describes the statistical weights of states with different energies and particle numbers.

And the Grand Potential:

$$\Omega = -k_B T \ln \mathcal{Z}$$

The grand potential is a thermodynamic potential for systems that can exchange both energy and particles with a reservoir. It is directly linked to the grand canonical partition function, with the logarithm describing the weighted contributions of all possible states.

Works Cited

Electromagnetic Theory & Computation: A Topological Approach — Gross & Kotiuga. Accessed December 24, 2025. https://library.slmath.org/books/Book48/files/gross-kotiuga.pdf

Notes on Linear Algebra & Functional Analysis — B. Bamieh. Accessed December 24, 2025. https://sites.me.ucsb.edu/~bamieh/pubs_new/MM.pdf

Introduction to Functional Analysis — D. Daners. Accessed December 25, 2025. https://www.maths.usyd.edu.au/u/athomas/FunctionalAnalysis/daners-functional-analysis-2017.pdf

Lie Algebras and Applications — F. Lachello. Accessed January 1st, 2026. http://ndl.ethernet.edu.et/bitstream/123456789/73867/1/87.pdf.pdf

Calculus of Variations — R. Cristoferi & S. Sonner. Accessed January 1st, 2026. https://www.math.ru.nl/~ssonner/CoV_lecturenotes(2022_23).pdf

The Principles of Quantum Mechanics — P. A. M. Dirac. Accessed December 26, 2025. https://gilles.montambaux.com/files/histoire-physique/Dirac-Principles%20of%20Quantum%20Mechanics%20.pdf

A New Approach to Quantum Theory — R. P. Feynman. Accessed December 26, 2025. http://files.untiredwithloving.org/thesis.pdf

Quantum Electrodynamics (Frontiers in Physics) — R. P. Feynman. Accessed December 24, 2025. https://ia802906.us.archive.org/21/items/in.ernet.dli.2015.176864/2015.176864.Quantum-Electrodynamics.pdf

Solid-State Theory: An Introduction — U. Rossler. Accessed December 24, 2025. https://nzdr.ru/data/media/biblio/kolxoz/P/PS/Roessler%20U.%20Solid%20State%20Theory..%20An%20Introduction%20(2ed.,%20Springer,%202009)(ISBN%203540927611)(O)(409s)_PS_.pdf

Introduction to Plasma Physics and Controlled Fusion — F. F. Chen. Accessed December 24, 2025. https://www.seas.ucla.edu/~ffchen/Archive/Chen020.pdf

Mathematical Introduction to Quantum Electrodynamics — J. Derezinski. Accessed January 10, 2026. https://www.fuw.edu.pl/~derezins/qed-slides-ihp.pdf

Evasion Paths in Mobile Sensor Networks — H. Adams. Accessed January 12th, 2026. https://people.clas.ufl.edu/henry-adams/files/Thesis_HenryAdams.pdf

Introduction to Quantum Mechanics- D. J. Griffiths & D. F. Schroeter. Accessed January 12th, 2026. https://eclass.uoa.gr/modules/document/file.php/CHEM248/Griffiths%20-%20Introduction%20to%20Quantum%20Mechanics%203rd%20ed%202018.pdf

Introduction to Elementary Particle Physics — D. J. Griffiths. Accessed February 1st, 2026. https://ia600502.us.archive.org/15/items/GriffithsD.IntroductiontoelementaryparticlesWiley1987/GriffithsD.Introductiontoelementaryparticles%28Wiley%2C1987%29.pdf

Quantum Magnetohydrodynamics — F. Haas. Accessed February 2nd, 2026. https://arxiv.org/pdf/physics/0503021

Nuclear Physics from Quantum Chromodynamics — S. R. Beane, W. Detmold, K. Originos, M.J. Savage. Accessed February 3rd, 2026. https://arxiv.org/pdf/1004.2935

Quantum Statistical Mechanics — M. Schwartz. Accessed March 2nd, 2026. https://schwartz.scholars.harvard.edu/sites/g/files/omnuum7046/files/schwartz/files/10-quantumstatmech.pdf

Spatio-Temporal Analysis of Urban Air Pollution Dynamics in Delhi-NCR: Hotspot Identification, Seasonal Trends, and Inter-Pollutant Correlations By Aadish Garg

Abstract

Delhi consistently ranks among the world's most polluted cities, with severe and spatially uneven air quality. This study examines the spatio-temporal dynamics of six criteria pollutants, namely PM2.5, PM10, NO2, ozone, SO2 and CO, across the National Capital Territory of Delhi for three representative years spanning 2021 to 2025, with the aim of identifying pollution hotspots and characterising seasonal trends and inter pollutant relationships. Daily concentration data from monitoring stations operated under the Central Control Room for Air Quality Management were cleaned and aggregated into monthly and seasonal means. High resolution seasonal concentration surfaces were generated for each pollutant using Inverse Distance Weighting interpolation within a Geographic Information Systems framework, and pairwise inter pollutant correlations were computed for each season to infer common sources. Results reveal two contrasting spatial behaviours: particulate matter and NO2 form persistent hotspots over the densely built up urban core and traffic corridors, whereas SO2 and CO remain low and comparatively uniform. Particulate matter, NO2 and CO peak in winter and the post monsoon season and fall sharply during the monsoon, while ozone shows the opposite cycle, confirming its secondary photochemical origin. Seasonally resolved correlations reveal strong particulate to combustion coupling during the high pollution seasons and a distinct NO2 to SO2 industrial signature, with ozone decoupled from all primary pollutants. The findings support geographically targeted, source specific control strategies and offer a transparent, reproducible workflow for air quality assessment in other rapidly urbanising cities.

Keywords Spatio-temporal analysis, air quality monitoring, pollution hotspots, Geographic Information Systems (GIS), urban air pollution.

Introduction

Ambient air pollution has become one of the foremost environmental threats to human health worldwide, and nowhere is the crisis more acute than in India (Balakrishnan et al.; Gulia et al.). Long-term exposure to fine particulate matter ($PM_{2.5}$) has been linked to an estimated 3.8 million deaths in India between 2009 and 2019 under the national air-quality standard, equivalent to roughly 5% of all mortality, a figure that rises to about 1.5 million deaths annually when benchmarked against World Health Organization (WHO) guideline concentrations (Jaganathan et al.; Agarwal & Kumar). The Global Burden of Disease Study similarly attributed close to one million deaths in India to ambient air pollution in a single year (Pandey et al.). Within this national context, the Delhi–National Capital Region (Delhi-NCR) is repeatedly identified as one of the most polluted urban agglomerations on Earth; Delhi has ranked as the world's most polluted capital city for several consecutive years, recording annual mean $PM_{2.5}$ concentrations exceeding 80 µg/m³ which is more than sixteen times the WHO annual guideline (IQAir).

The severity of Delhi's air quality arises from a dense concentration of emission sources superimposed on unfavourable geography and meteorology (Ansari et al.). Rapid urban population growth, industrialisation, a large and expanding vehicle fleet, extensive construction activity, and domestic and commercial combustion together generate high baseline emissions across the territory (Thanvisitthpon et al; Gorai et al). Delhi's landlocked position within the Indo-Gangetic Plain deprives it of the ventilating coastal breezes that disperse pollutants in maritime cities, so emissions accumulate under stagnant conditions. This effect is strongly seasonal: during winter, low mixing heights, frequent temperature inversions, and calm winds trap pollutants near the surface, producing repeated "severe" air-quality episodes, whereas the summer monsoon temporarily reduces particulates and improves air quality (Gayen et al.). Consequently, pollutant concentrations in Delhi are not only high but also highly variable in both space and time, a characteristic that any meaningful assessment must explicitly capture.

A defining feature of the post-monsoon deterioration is the influence of large-scale agricultural crop-residue burning in the neighbouring states of Punjab and Haryana (Singh et al.; Rahman et al.). Following the Kharif harvest in October and November, the open burning of paddy straw releases enormous quantities of particulate matter, which prevailing north-westerly winds advect into Delhi (Nagar et al.). Mean ambient $PM_{2.5}$ concentrations at Delhi monitoring stations have been observed to climb sharply through the monsoon, reaching monthly averages above 230 $\mu g/m^3$ in November (Rahman et al.), and modelling studies attribute a substantial share of the post-monsoon particulate load to this transported regional source (Kulkarni et al.). The interplay between locally generated emissions and this external, seasonally pulsed contribution makes Delhi-NCR a particularly complex system in which to disentangle pollution sources.

Air quality in the region is governed by six criteria pollutants: $PM_{2.5}$, coarse particulate matter (PM_{10}), nitrogen dioxide (NO_2), ground-level ozone (O_3), sulphur dioxide (SO_2), and carbon monoxide (CO); that differ markedly in their sources, chemistry, and spatial behaviour. Primary combustion products such as PM, NO_2, and CO tend to be co-emitted from traffic and industry (Awasthi et al.; Kumar et al.), whereas O_3 is a secondary pollutant formed photochemically and can therefore behave inversely to its precursors. These pollutants are monitored through a network of Continuous Ambient Air Quality Monitoring Stations, but such stations provide measurements only at discrete points, leaving the concentration field between them unknown. To convert sparse point observations into the continuous spatial surfaces required for hotspot identification, spatial interpolation techniques are widely employed. Inverse Distance Weighting (IDW), a deterministic method that estimates unsampled locations as a distance-weighted average of surrounding observations, is among the most common and computationally efficient approaches used in Geographic Information Systems (GIS) for air-quality mapping (Ismain et al.; Shepard).

Despite a substantial body of research on Delhi's air pollution, important gaps remain. Many studies concentrate on a single pollutant; most often $PM_{2.5}$; or on a single season or short episode, and comparatively few simultaneously map multiple pollutants, characterise their

seasonal behaviour, and quantify the statistical relationships among them over several years. Yet it is precisely the combination of these perspectives that is most informative: spatial mapping reveals where pollution concentrates, seasonal analysis reveals when and why concentrations change, and inter-pollutant correlation analysis provides a quantitative basis for inferring shared emission sources. Strong positive correlations among combustion-related pollutants point to common vehicular and industrial origins, whereas weak or negative correlations as typically seen between O_3 and its precursors, reflect distinct photochemical processes rather than co-emission. Higher $PM_{2.5}/PM_{10}$ ratios in winter and post-monsoon further signal shifts in the dominant particulate sources across seasons. Integrating these analytical dimensions over a multi-year window therefore offers a more complete diagnostic picture than any single approach alone.

The present study addresses this gap through an integrated spatio-temporal analysis of the six criteria pollutants across Delhi-NCR for three representative years: 2021, 2023, and 2025, thus spanning a five-year period. Daily observations from the region's monitoring stations are cleaned and aggregated to monthly and seasonal means, high-resolution seasonal concentration surfaces are generated for each pollutant using IDW interpolation within a GIS framework, and pairwise inter-pollutant correlations are computed to infer common sources. The objectives are threefold: to identify persistent pollution hotspots and relate them to underlying land use, to characterise the seasonal and inter-annual variability of each pollutant, and to quantify inter-pollutant relationships as a basis for source attribution. By combining multi-pollutant GIS mapping, seasonal trend analysis, and correlation analysis within a single, reproducible workflow, this study aims to support geographically targeted emission-control strategies and to offer a transferable methodological template for air-quality assessment in other rapidly urbanising cities.

Study Area

The study area comprises the National Capital Territory (NCT) of Delhi, in northern India within the Indo-Gangetic Plain, together with its bordering districts in the wider Delhi–National Capital Region (Delhi-NCR) (Figure 1). The NCT extends over approximately 1,483 km² between latitudes 28.40°N and 28.88°N and longitudes 76.84°E and 77.35°E, bordered by Haryana to the west, south, and north and by Uttar Pradesh to the east. The terrain is largely flat, forming part of the alluvial plain of the Yamuna River, with the low Aravalli ridge crossing the south-western portion of the city. Its inland, landlocked setting deprives the region of the ventilating coastal breezes that disperse pollutants in maritime cities, so emissions accumulate under stagnant conditions (Gorai et al.).

As shown in Figure 1(b), built-up land dominates the central urban core (~42%), surrounded by a peri-urban cropland belt (~31%) and patches of tree cover (~19%), with the Yamuna running north to south. The NCT alone houses well over 16 million residents and the broader NCR exceeds 30 million, sustaining intense vehicular traffic, industrial activity, construction, and energy demand. The region experiences a semi-arid, monsoon-influenced subtropical climate with hot summers, a short but intense south-west monsoon, and cold, foggy winters. This pronounced seasonality strongly modulates air quality: winter brings low mixing heights, temperature inversions, and calm winds that trap pollutants near the surface, whereas

monsoon rainfall temporarily scavenges particulates (Gayen et al.). The post-monsoon period coincides with large-scale crop-residue burning in neighbouring Punjab and Haryana, whose smoke is advected into the city by north-westerly winds and sharply degrades air quality (Nagar et al.; Kulkarni et al.).

These conditions have made Delhi a persistent global pollution hotspot; it remains the world's most polluted capital, with a recent annual mean $PM_{2.5}$ concentration exceeding 80 µg/m³ — over sixteen times the WHO guideline (IQAir, 2025). The combination of severe, spatially heterogeneous pollution, diverse overlapping emission sources, and strong meteorological control makes Delhi-NCR an ideal setting for spatio-temporal analysis. Air quality across the region is monitored by a network of Continuous Ambient Air Quality Monitoring Stations (CAAQMS) operated by the Central Pollution Control Board (CPCB), the Delhi Pollution Control Committee (DPCC), and the Indian Institute of Tropical Meteorology (IITM).

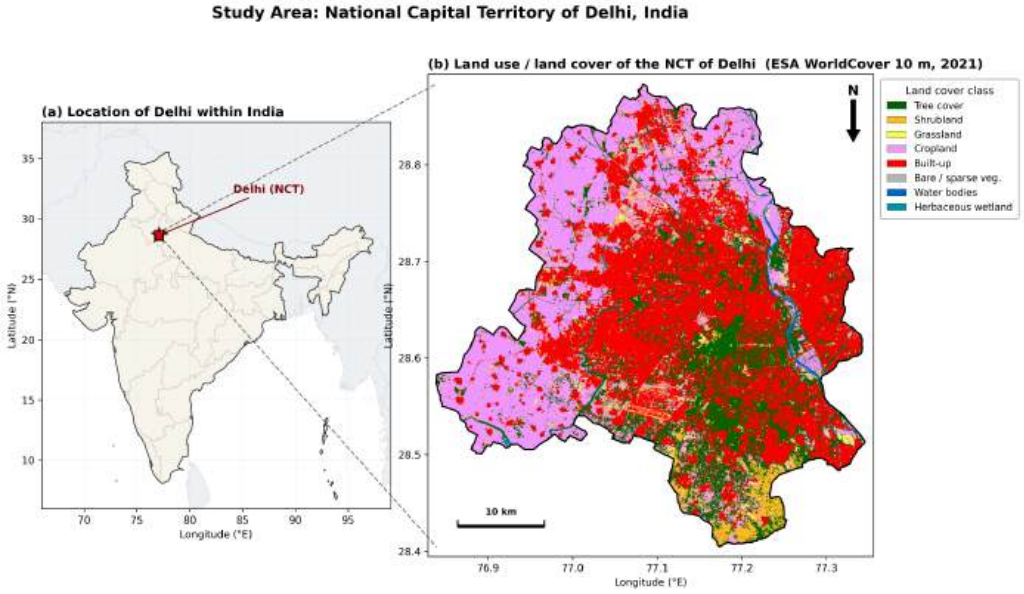


Figure 1: Study area. (a) Location of the National Capital Territory (NCT) of Delhi within India; (b) land use/land cover of the NCT of Delhi derived from ESA WorldCover (10 m, 2021), clipped to the administrative boundary.

Data and Methodology

Data Source and Description

Ground-level pollutant concentration data were sourced from the Central Control Room for Air Quality Management operated by the Central Pollution Control Board (CPCB), which aggregates continuous ambient air quality measurements reported by Continuous Ambient Air Quality Monitoring Stations (CAAQMS) operated by CPCB, the Delhi Pollution Control Committee (DPCC), and the India Meteorological Department / IITM network across the region. A total of 41 monitoring stations distributed across Delhi-NCR were included in the analysis.

Six criteria pollutants were examined: fine particulate matter ($PM_{2.5}$), coarse particulate matter (PM_{10}), nitrogen dioxide (NO_2), ground-level ozone (O_3), sulphur dioxide (SO_2), and carbon monoxide (CO). Particulate and gaseous pollutants ($PM_{2.5}$, PM_{10}, NO_2, O_3, SO_2) are expressed in micrograms per cubic metre (µg/m³), while CO is expressed in milligrams per cubic metre (mg/m³), consistent with CPCB reporting conventions and the National Ambient Air Quality Standards (NAAQS).

The analysis focuses on three representative years 2021, 2023, and 2025, spanning a five-year period (2021–2025). Sampling the study window at two-year intervals captures the medium-term evolution of spatial pollution patterns while keeping the cartographic output interpretable, and the same three-year set underpins the seasonal mapping, temporal comparison, and inter-pollutant correlation analyses.

Data Pre-processing and Quality Control

Raw station records were subjected to a sequence of cleaning steps prior to analysis. Station labels in the concentration files and in the coordinate reference file frequently differed in formatting (e.g., appended operating-agency suffixes such as "– DPCC", "– CPCB", or "– IITM", inconsistent use of underscores, commas, and full stops, and redundant inclusion of the word "Delhi"). A normalisation routine was applied to both datasets to strip agency suffixes and punctuation, remove the redundant place name, collapse repeated whitespace, and convert to a common lower-case key. Concentration records were then merged with station coordinates (latitude, longitude) on this harmonised key, and a completeness check confirmed that every concentration record was successfully matched to a geographic location before the analysis proceeded.

Daily concentration measurements were aggregated to monthly mean values at each station. Monthly aggregation reduces the influence of short-term measurement noise and transient episodes, yields a consistent temporal resolution across all stations, and provides the basis for both the seasonal spatial maps and the long-term trend series. Missing monthly values were excluded on a pairwise basis rather than imputed. When computing a seasonal mean for a given station, the average was taken only over the months with valid data (missing months were skipped), and stations lacking any valid observation for a particular season were omitted from the interpolation for that season only. This preserves the maximum available information at each station without introducing artificial values.

Seasonal Classification

To capture the strong meteorological modulation of air quality in Delhi-NCR, each calendar year was partitioned into four seasons following the regional climatic cycle:

Table 1. Classification of the four seasons used in the analysis and their constituent months, following the regional climatic cycle of Delhi.

Season	Constituent months
Winter	December, January, February
Summer / Pre-Monsoon	March, April, May, June
South-West Monsoon	June, July, August, September
Post-Monsoon	October, November

For each station, year, and season, a seasonal mean concentration was computed as the average of the constituent monthly means (skipping missing months as described above). These seasonal means constitute the input values for spatial interpolation. **Methodological note:** In the current season definitions June is shared between the Summer/Pre-Monsoon and South-West Monsoon windows. This reflects the transitional onset of the monsoon but produces partial overlap between the two seasons.

Spatial Interpolation and GIS Mapping

Continuous spatial surfaces of seasonal mean concentration were generated from the discrete station values using Inverse Distance Weighting (IDW), a deterministic interpolation method in which the estimated value at any unsampled location is a distance-weighted average of the surrounding observations. The interpolated value at a grid location was computed as per equations below.

$$\hat{Z}(s_0) = \frac{\sum_{i=1}^{n} w_i Z(s_i)}{\sum_{i=1}^{n} w_i}, \qquad w_i = \frac{1}{d_i^p}$$

with d_i the Euclidean distance (in decimal degrees) from grid location s_0 to monitoring station s_i, $Z(s_i)$ the observed seasonal mean at that station, and p the power (distance-decay) parameter. A power of **p = 2** was adopted, giving greater influence to nearby stations — an appropriate and widely used choice at the intra-urban scale, where the station network is dense relative to the domain extent. A negligible offset was applied to zero distances to avoid singularities at station locations.

Interpolation was performed on a regular **300 × 300** grid spanning the bounding box of the NCT boundary, yielding a spatial resolution of the order of a few hundred metres across the territory. To restrict estimates to the administrative domain, every grid cell was tested for containment within the NCT polygon, and cells falling outside the boundary were masked (set to no-data) so that interpolated surfaces were clipped exactly to the shape of Delhi.

The resulting surfaces were rendered as filled contour maps in a GIS-style layout. To correct for latitude-dependent longitudinal compression, the map aspect ratio was scaled by 1 / cos(latitude) at the domain mean latitude, giving geographically faithful proportions. Each map

was annotated with a north arrow, a 10 km scale bar (with the degree-to-kilometre conversion evaluated at the domain mean latitude), station locations overlaid as point markers, and the administrative boundary outline. For every pollutant, all seasons and all three mapped years shared a single common colour scale, with lower and upper limits fixed to the global minimum and maximum across the entire set of seasonal means. This deliberate design choice ensures that colours are directly comparable across seasons and years, so that apparent differences in the maps reflect genuine differences in concentration rather than independent per-panel rescaling. A red–yellow–green diverging colour ramp (reversed, such that green denotes lower and red denotes higher concentrations) was used for intuitive interpretation, and a 16-level contour classification was applied uniformly. The final products are 2 × 2 seasonal composite figures, one per pollutant per representative year.

Inter-Pollutant Correlation Analysis

To investigate shared emission sources and photochemical relationships among pollutants, pairwise correlations were computed across the co-located, time-matched monthly concentration series for all six pollutants. Correlation coefficients were calculated both across the full dataset and stratified by season, and the results were summarised as a correlation matrix and visualised as a heat map. Pearson (linear) correlation coefficients were used in this study. Strong positive correlations among combustion-related pollutants (e.g., $PM_{2.5}$, PM_{10}, CO, and NO_2) were interpreted as evidence of common vehicular and combustion sources, while the characteristically inverse or lagged relationship between NO_2 and O_3 was interpreted in the context of photochemical ozone formation and titration.

Results and Discussion

Spatial Distribution and Pollution Hotspots

The interpolated seasonal surfaces reveal that the six pollutants differ strongly in how their concentrations are distributed across the territory. Particulate matter is the most spatially extensive problem. For PM2.5 (Figure 2) and PM10 (Figure 3), the highest concentrations form a broad zone across the northern and north central districts that extends into the central urban core, most clearly during winter and the post monsoon season, when almost the entire territory shifts to orange and deep red on the shared colour scale. A recurring pocket of comparatively lower particulate loading appears over the southwestern rural fringe near 28.55 N, 76.95 E, visible as a persistent green area in nearly every panel. The close spatial agreement between the PM2.5 and PM10 fields indicates that the two size fractions share their dominant source areas across the city.

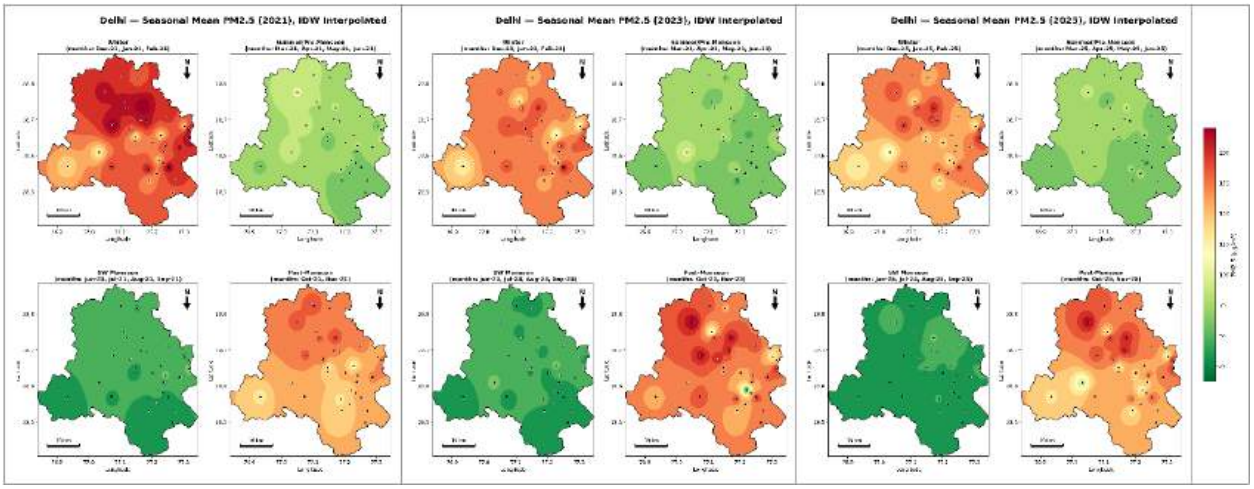


Figure 2. Seasonal mean PM2.5 concentrations (µg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation. Each block of four panels shows the winter, summer/pre monsoon, south west monsoon and post monsoon seasons for one year.

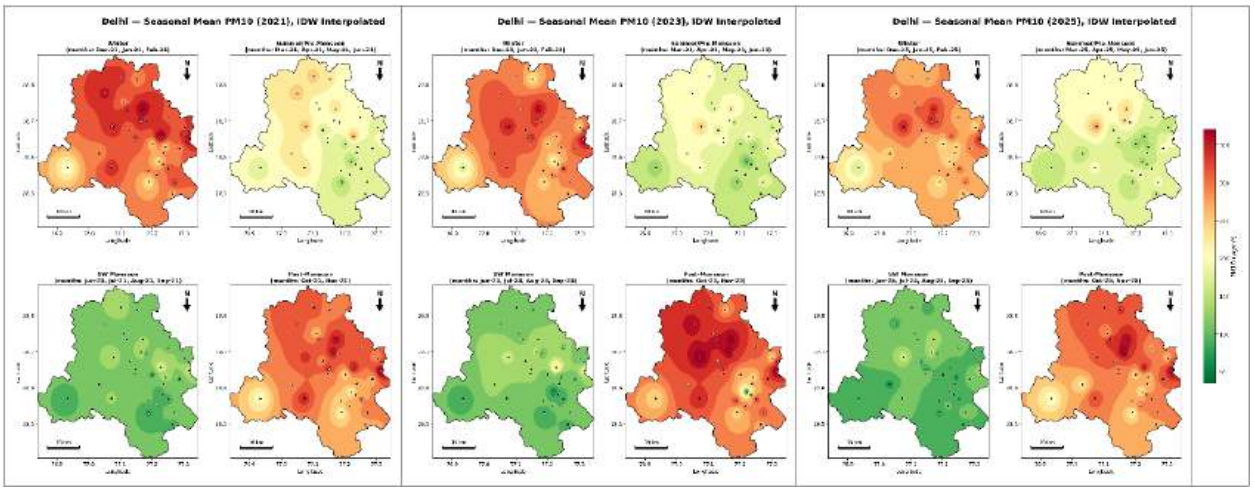


Figure 3. Seasonal mean PM10 concentrations (µg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation.

Nitrogen dioxide (Figure 4) is more sharply localised. Rather than a single broad plume, NO2 concentrates over the dense central core and along the eastern and southeastern margins of the city, including the trans Yamuna corridor, where distinct red hotspots recur across years and intensify markedly in the winter of 2025. These locations coincide with the most heavily trafficked and built up parts of Delhi shown in Figure 1(b), consistent with a vehicular and combustion origin. The western and outer rural areas remain green throughout.

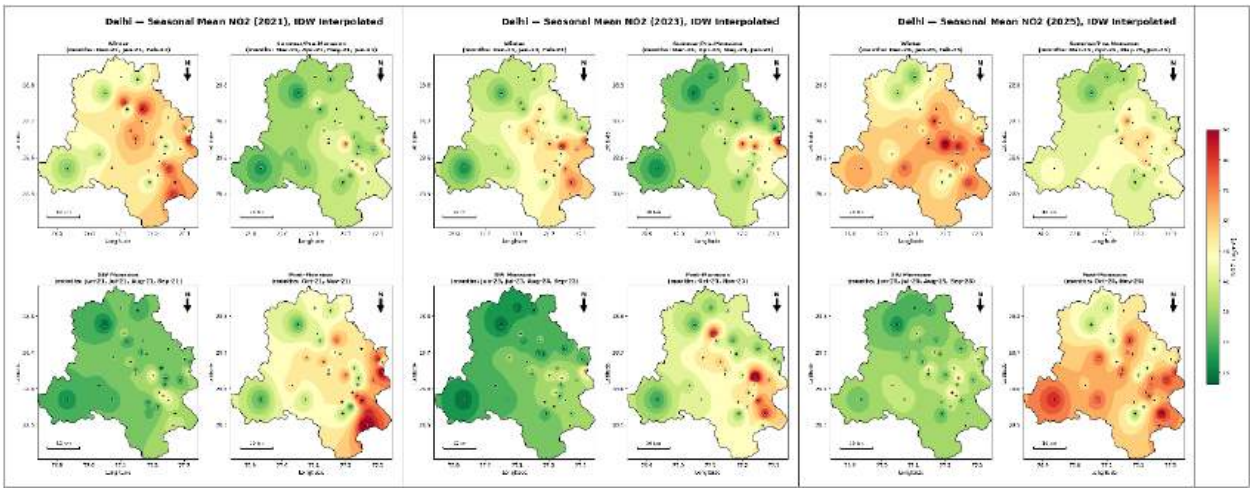


Figure 4. Seasonal mean NO2 concentrations (µg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation.

Ground level ozone (Figure 5) shows the opposite spatial tendency. Its highest values occur not in the congested core but over the southern and peripheral parts of the territory, particularly during the summer and pre monsoon period, with a pronounced hotspot in the south and southeast in 2021 and a strong western maximum in 2025. This inverse spatial relationship to NO2 reflects the photochemical nature of ozone and its suppression close to fresh emission sources, where high concentrations of nitric oxide consume ozone near the surface.

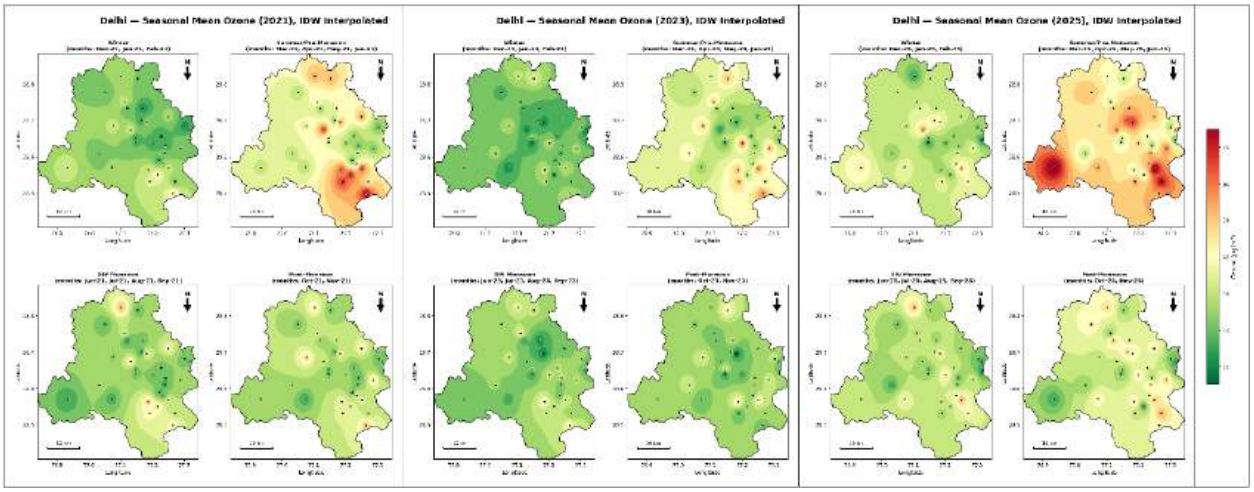


Figure 5. Seasonal mean ground level ozone (O3) concentrations (µg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation.

Sulphur dioxide (Figure 6) and carbon monoxide (Figure 7) are, by contrast, comparatively uniform and low across most of the territory. Both remain predominantly green over the study period, with elevated values confined to a small number of isolated point hotspots rather than broad zones. For SO2 these localised maxima become more prominent in 2025, appearing as sharp red points in the north central and southern city that likely mark specific

industrial or fuel combustion sources. CO displays moderate and largely even concentrations, with intermittent hotspots in the central and southeastern city during the high pollution seasons.

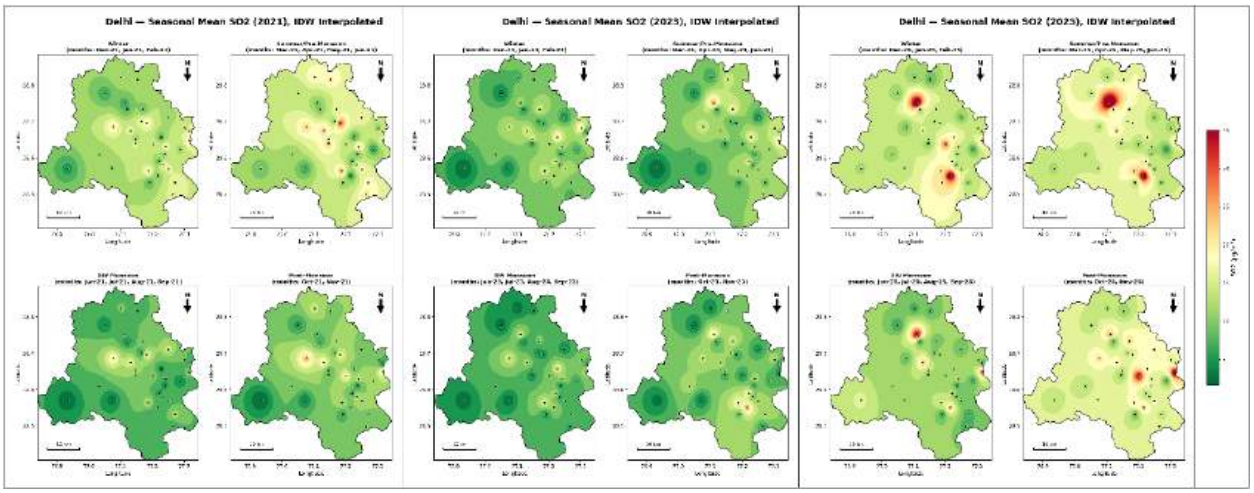


Figure 6. Seasonal mean SO2 concentrations (µg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation.

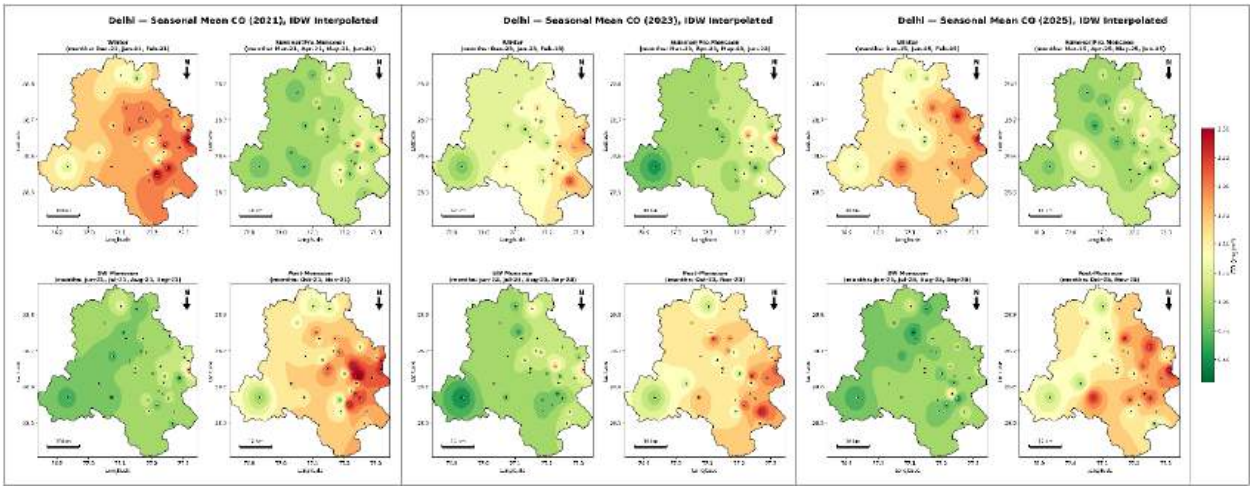


Figure 7. Seasonal mean CO concentrations (mg/m³) across the NCT of Delhi for 2021, 2023 and 2025, generated by IDW interpolation.

Taken together, the maps distinguish two contrasting behaviours. The spatially heterogeneous pollutants PM2.5, PM10 and NO2 form persistent hotspots aligned with the built up core and the major traffic corridors, whereas SO2 and CO are distributed more evenly at generally low levels. Because the monitoring network becomes sparser toward the peri urban boundary, interpolated values along the outer edges should be read with greater caution than those in the densely monitored centre, where station coverage constrains the interpolation most tightly.

Seasonal and Inter-Annual Variability

Read season by season, the maps expose a strong and highly consistent meteorological cycle that governs most of the pollutants. Particulate matter shows the clearest pattern. For PM2.5 (Figure 2) and PM10 (Figure 3), winter and the post monsoon season stand out as the two peak periods, when large parts of the territory move into the orange and red bands of the shared colour scale, with winter PM2.5 approaching the upper end of the range across the northern and central districts. Concentrations then fall sharply through the summer and reach their annual minimum during the south west monsoon, when almost the entire territory turns green and PM2.5 drops to the lowest classes on the scale. The amplitude of this cycle is large, representing a severalfold decrease from the winter maximum to the monsoon minimum, and it is driven by the seasonal contrast between the shallow, stable winter boundary layer, which traps emissions near the surface, and the deep mixing and wet scavenging that accompany monsoon rainfall.

Carbon monoxide (Figure 7) follows the same rhythm as particulate matter, with the highest values in winter, a clear decline through summer and the monsoon, and a secondary rise in the post monsoon season. This shared timing points to common combustion sources and to the same meteorological control acting on primary pollutants. Nitrogen dioxide (Figure 4) also peaks in the cooler months and reaches its lowest and most spatially restricted values during the monsoon, when only isolated green cores remain across the western and northern districts.

Ozone behaves in the opposite way and provides the most instructive contrast in the dataset. As Figure 5 shows, ozone concentrations are highest during the summer and pre monsoon period, when strong solar radiation and elevated temperatures promote photochemical production, and are lowest in winter and the monsoon. This inverse seasonal cycle relative to the primary pollutants confirms that ozone is a secondary product of atmospheric chemistry rather than a directly emitted species. Sulphur dioxide (Figure 6) is the exception to the seasonal rule, showing only a weak and irregular cycle; its concentrations remain low across all four seasons and are dominated instead by a small number of localised point sources.

The post monsoon intensification of PM2.5 and PM10 deserves particular emphasis. In this season the particulate fields not only rise across the whole territory but also develop concentrated red cores in the central and northern city, most strongly in 2023 and 2025. This monsoon deterioration coincides with the period of large scale crop residue burning in the neighbouring agricultural states, whose smoke is transported into the city under prevailing north westerly winds and superimposed on already accumulating local emissions.

Comparison across the three sampled years should be read with caution, since 2021, 2023 and 2025 are not consecutive and each season reflects the particular meteorology of that year. With that caveat, a few changes are visible. Winter NO2 appears more widespread and intense in 2025 than in the earlier years, and sulphur dioxide develops sharper and more numerous point hotspots in 2025 that are largely absent in 2021 and 2023. These features are best described as differences between the sampled years rather than as established long term trends.

Inter-Pollutant Correlations and Source Attribution

The seasonal correlation heatmaps (Figure 8) quantify how the six pollutants co vary and provide a statistical basis for inferring shared sources. The single most consistent relationship across every season is the very strong positive correlation between PM2.5 and PM10, which ranges from about 0.68 during the monsoon to 0.90 in the post monsoon season. This tight coupling confirms that the two particulate fractions are largely co emitted and respond together to the same meteorological controls, echoing the close spatial agreement already seen in their maps.

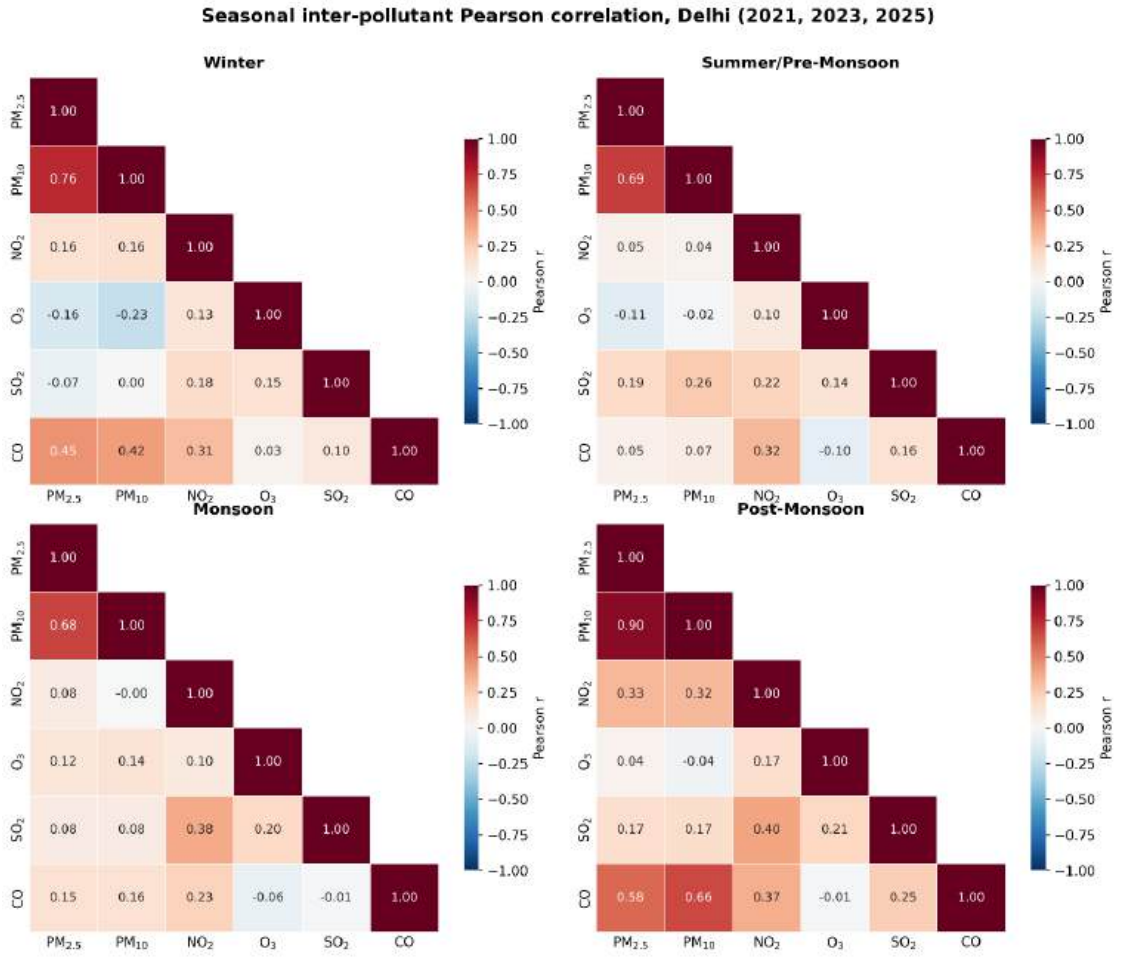


Figure 8. Seasonal inter pollutant Pearson correlation matrices for the six criteria pollutants across the NCT of Delhi (2021, 2023 and 2025), shown separately for winter, summer/pre monsoon, monsoon and post monsoon. Cell values are Pearson correlation coefficients (r), with red indicating positive and blue negative associations; only the lower triangle is shown to avoid duplication.

The behaviour of the particulate pollutants with carbon monoxide is the most revealing feature of the analysis, because it changes markedly with season. Coupling is tightest in the post monsoon period, when CO correlates with PM2.5 at about 0.58 and with PM10 at about 0.66, and remains moderate in winter, at roughly 0.45 and 0.42 respectively. By contrast, these correlations collapse almost to zero during summer and the monsoon. The strong particulate to

CO association in the post monsoon and winter months, precisely when concentrations peak, is a clear combustion signature and is consistent with the shared contribution of vehicular emissions and seasonal biomass burning to both particulate matter and carbon monoxide during the high pollution period.

Nitrogen dioxide forms a secondary cluster that points more toward mixed combustion and industrial activity than to the particulate sources. Its correlation with PM2.5 and PM10 is generally weak to moderate and strongest in the post monsoon season, at about 0.33 and 0.32. More notably, NO2 shows its most stable relationship with sulphur dioxide, correlating at around 0.38 to 0.40 in the monsoon and post monsoon seasons. This persistent NO2 to SO2 link, distinct from the particulate cluster, is characteristic of fixed fuel combustion and industrial emissions rather than the dispersed vehicular and burning sources that dominate particulate loading.

Ozone stands apart from every other pollutant, confirming the interpretation drawn from its maps. In each season its correlations with PM2.5, PM10 and CO are near zero or slightly negative, including values such as negative 0.16 with PM2.5 and negative 0.23 with PM10 in winter, and negative 0.10 with CO in summer. The absence of any positive association with the primary combustion pollutants demonstrates that ozone is not co emitted but is instead formed by photochemical processes, and the weak negative signs are consistent with its suppression near fresh sources of nitrogen oxides. The only pollutant with which ozone shows a mild recurring positive relationship is sulphur dioxide, at around 0.15 to 0.21, which most plausibly reflects a shared response to the warm, sunlit conditions of the pre monsoon season rather than a common source.

The seasonal stratification is what gives these results their value. A single correlation computed over the whole dataset would blend these contrasting regimes and overstate the apparent coupling, because most pollutants happen to rise together in winter for meteorological rather than source related reasons. By separating the seasons, the analysis shows that the particulate to combustion linkage is genuinely strongest in the post monsoon and winter periods, weakens under the well mixed and rain washed conditions of the monsoon, and never extends to ozone. Two cautions apply to the interpretation. Correlation cannot establish causation, and the coefficients reflect statistical association only, so the source inferences drawn here should be treated as consistent with, rather than proof of, the proposed emission origins. Nonetheless, the convergence of three independent lines of evidence, namely the spatial hotspots, the seasonal cycles, and these inter pollutant correlations, points coherently toward combustion and vehicular activity concentrated in the urban core, amplified by winter meteorology and by the transported post monsoon burning signal, as the principal drivers of Delhi's air quality.

Conclusion

This study examined the spatio-temporal dynamics of six criteria pollutants, namely PM2.5, PM10, NO2, ozone, SO2 and CO, across the National Capital Territory of Delhi for three representative years spanning 2021 to 2025, combining IDW based seasonal mapping within a GIS framework with inter pollutant correlation analysis. Three consistent findings emerge. First,

the pollutants divide into two spatial groups. Particulate matter and nitrogen dioxide are strongly heterogeneous and form persistent hotspots over the densely built up northern and central districts and the eastern traffic corridors, aligned with the urban land cover of the study area. Sulphur dioxide and carbon monoxide remain comparatively low and uniform, punctuated only by isolated point sources that became more pronounced for SO2 in 2025, while ozone peaks over the southern and peripheral areas rather than the congested core.

Second, air quality is governed by a strong seasonal cycle. Particulate matter, NO2 and CO reach their maxima in winter and the post monsoon season, driven by a shallow, stable boundary layer and, in monsoon, by transported crop residue burning, and fall to their minima during the monsoon through enhanced mixing and wet scavenging. Ozone alone shows the opposite cycle, peaking in the summer and pre monsoon months, confirming its secondary, photochemical origin. Third, the seasonally resolved correlations reveal shared sources. The very strong PM2.5 to PM10 association, together with the strong particulate to CO coupling that emerges in the post monsoon and winter periods, points to combustion and vehicular activity as the dominant contributors during the high pollution seasons, while the separate NO2 to SO2 relationship reflects a fixed combustion and industrial signature. Ozone remains statistically decoupled from all primary pollutants.

Together, these three lines of evidence point consistently toward combustion and vehicular emissions concentrated in the urban core, amplified by winter meteorology and the transported post monsoon burning signal, as the principal drivers of Delhi's air quality. The stability of the particulate and NO2 hotspots argues for geographically targeted, source specific interventions rather than uniform measures, and the monsoon peak underscores the need for coordinated regional action on agricultural burning. The main limitations are the use of three non consecutive years, the reduced reliability of interpolation near the sparsely monitored margins, and the correlational nature of the source inferences. Even so, the integrated workflow offers a transparent, reproducible template for air quality assessment transferable to other rapidly urbanising cities.

Works Cited

Agarwal, V., and A. Kumar. “The Changes in the Air Quality of Wazirpur, Delhi Due to the COVID-19 Shutdown.” Cleaner Chemical Engineering, vol. 1, no. 2, 2022, article 100001. https://doi.org/10.1016/j.clce.2022.100001.

Ansari, N., et al. “Seasonal Patterns of Air Pollution in Delhi: Interplay between Meteorological Conditions and Emission Sources.” Environmental Geochemistry and Health, vol. 47, no. 5, 2025, article 175. https://doi.org/10.1007/s10653-025-02474-0.

Awasthi, A., et al. “Biomass-Burning Sources Control Ambient Particulate Matter, but Traffic and Industrial Sources Control Volatile Organic Compound (VOC) Emissions and Secondary-Pollutant Formation during Extreme Pollution Events in Delhi.” Atmospheric Chemistry and Physics, vol. 24, no. 18, 2024, pp. 10279–10304. https://doi.org/10.5194/acp-24-10279-2024.

Balakrishnan, K., et al. “The Impact of Air Pollution on Deaths, Disease Burden, and Life Expectancy across the States of India: The Global Burden of Disease Study 2017.” The Lancet Planetary Health, vol. 3, no. 1, 2019, pp. e26–e39. https://doi.org/10.1016/S2542-5196(18)30261-4.

Gayen, B. K., et al. “Exploring the Effect of Waterbodies Coupled with Other Environmental Parameters to Model PM2.5 over Delhi-NCT in Northwest India.” Atmospheric Pollution Research, vol. 13, no. 12, 2022, article 101614. https://doi.org/10.1016/j.apr.2022.101614.

Gorai, A. K., et al. “Spatio-Temporal Variation of Particulate Matter (PM2.5) Concentrations and Its Health Impacts in a Mega City, Delhi in India.” Environmental Health Insights, vol. 12, 2018, pp. 1–10. https://doi.org/10.1177/1178630218792861.

Gulia, S., et al. “Evolution of Air Pollution Management Policies and Related Research in India.” Environmental Challenges, vol. 6, 2022, article 100431. https://doi.org/10.1016/j.envc.2021.100431.

IQAir. World Air Quality Report 2024. IQAir, 2025.

Ismain, S. H., et al. “Spatial Distribution of Particulate Matter (PM2.5) in Klang Valley Using Inverse Distance Weighting Interpolation Model.” IOP Conference Series: Earth and Environmental Science, vol. 1217, 2023, article 012033. https://doi.org/10.1088/1755-1315/1217/1/012033.

Jaganathan, S., et al. “Estimating the Effect of Annual PM2.5 Exposure on Mortality in India: A Difference-in-Differences Approach.” The Lancet Planetary Health, vol. 8, no. 12, 2024. https://doi.org/10.1016/S2542-5196(24)00248-1.

Kulkarni, S. H., et al. “How Much Does Large-Scale Crop Residue Burning Affect the Air Quality in Delhi?” Environmental Science & Technology, vol. 54, no. 8, 2020, pp. 4790–4799. https://doi.org/10.1021/acs.est.9b07038.

Kumar, V., P. C. S. Devara, and V. K. Soni. “Spatiotemporal Variations in Black Carbon Aerosol Concentration and Its Correlation with Meteorological Parameters and Other Pollutants in India.” MAUSAM, vol. 77, no. 1, 2026, pp. 219–234. https://doi.org/10.54302/cj91dh35.

Nagar, P. K., M. Sharma, and D. Das. “A New Method for Trend Analyses in PM10 and Impact of Crop Residue Burning in Delhi, Kanpur and Jaipur, India.” Urban Climate, vol. 27, 2019, pp. 193–203. https://doi.org/10.1016/j.uclim.2018.12.003.

Pandey, A., et al. “Health and Economic Impact of Air Pollution in the States of India: The Global Burden of Disease Study 2019.” The Lancet Planetary Health, vol. 5, no. 1, 2021, pp. e25–e38. https://doi.org/10.1016/S2542-5196(20)30298-9.

Rahman, M. H., N. Singh, S. Kundu, and A. Datta. “Potential Areas of Crop Residue Burning Contributing to Hazardous Air Pollution in Delhi during the Post-Monsoon Season.” Journal of Environmental Quality, vol. 51, no. 2, 2022, pp. 181–192. https://doi.org/10.1002/jeq2.20326.

Shepard, D. “A Two-Dimensional Interpolation Function for Irregularly-Spaced Data.” Proceedings of the 23rd ACM National Conference, 1968, pp. 517–524. https://doi.org/10.1145/800186.810616.

Singh, A., et al. “Impact of Crop Residue Burning on Regional Air Quality over the Indo-Gangetic Plain: A Comparative Study during Different Cultivation Patterns.” Scientific Reports, 2026.

Thanvisitthpon, N., K. Kallawicha, and H. J. Chao. “Effects of Urbanization and Industrialization on Air Quality.” Health and Environmental Effects of Ambient Air Pollution, edited by M. H. Dehghani et al., R. R. Karri, T. Vera, and S. K. M. Hassan, Academic Press, 2024, pp. 231–255. https://doi.org/10.1016/B978-0-443-16088-2.00003-X.

A Review of Environmental Changes Affecting Coastal Winds Near Nantucket and Effects on Sailing By Sammy Laliberte

Abstract

Dating as far back as 1000 BCE, sailors have depended on sea conditions and predictable wind patterns for safe navigation, shipping, recreation, and sport. However, as climate change worsens and changes atmospheric circulation, ocean temperatures, and storm behavior, research indicates coastal wind patterns and sea conditions are affected, creating new challenges for sailors worldwide. This paper aims to review existing literature on the short- and long-term effects of environmental changes on wind patterns and speeds along the coast of Nantucket, and how these changes affect sailing. The literature review will include research found through Google Scholar using combinations of search terms including "climate change" AND "wind speed" AND “sailing” AND "coastal wind patterns" AND "Nantucket" AND "wind climatology" AND "weather change." This review will show that changes in atmospheric circulation, ocean temperatures, and storm frequency change both wind and sea conditions. Although research varys depending on region, overall the literature suggests that more sporadic weather and wind conditions are becoming more frequent around the world and are a significant problem for sailors. It is critical to understand how the environment and anthropogenic factors influence coastal wind and weather because sailing safety, strategy, and marine transportation depend on reliable and safe conditions.

Introduction

Sailors have depended on sea conditions and predictable wind patterns for safe navigation, shipping, recreation, and sport. As sailors developed more sophisticated methods that helped the exchange of goods, ideas, and knowledge, marine travel began to expand all across the world (Schotte, 2019). Because sailing vessels rely on wind for propulsion, it is crucial to understand local weather patterns to ensure safe and effective travel (Anderson, 2003). It has also been found that the uneven heating of earths surface causes air to move, which generates wind and sea breezes. Anderson demonstrates how important sea breezes are to sailors, exclaiming, "If there is no other source of wind on a hot, summer day, a sailor can usually find at least a gentle shore breeze in the mid to late afternoon" (Anderson, 2003, p. 118). Even modern navigation in sailing still depends on weather patterns and accurate forecasts of wind and sea conditions. Currently, weather routing uses forecasts of waves and wind to help build a safe and efficient route (Ding & Zis, 2022). Therefore, understanding changes in wind patterns has become increasingly important for commercial and recreational sailors as the conditions they rely on continue to change.

Climate change has impacts on many of the atmospheric and marine processes that impact wind and sea conditions. For example, atmospheric circulation and land-use change have led to global shifts in wind patterns (Wu et al., 2017). Further, global wind speeds and wave heights have generally increased in recent years, with the biggest increases in the Southern

Climate change has impacts on many of the atmospheric and marine processes that impact wind and sea conditions. For example, atmospheric circulation and land-use change have led to global shifts in wind patterns (Wu et al., 2017). Further, global wind speeds and wave heights have generally increased in recent years, with the biggest increases in the Southern hemisphere (Zheng & Li, 2022). This indicates that climate change is impacting many of the conditions sailors rely on for safe navigation. Even though these studies showed global trends, climate change has different effects across different regions. For example, the Gulf of Maine has experienced warmer waters and different coastal conditions (Maine Climate Council [MCC] Scientific and Technical Subcommittee [STS], 2024).

While global trends are important to see climate related changes, regional studies are needed because coastal environments are affected differently by the changing atmospheric and ocean processes. For example, the Gulf of Maine and water around Nantucket are an important region to study changes to the ocean environment because they are surrounded by unique costal geography, local wind patterns, and circulation (Brown, 1998; MCC STS, 2024). Archer et al. (2016) found that the marine boundary layer over Nantucket Sound is unstable because the air near the ocean surface is warmer than the air above it, causing vertical mixing, turbulence, and air currents. These conditions have a large impact on offshore wind and make the region a good location for observing coastal wind patterns. Further, Manwell and Rogers (2002) found that the waters around Nantucket Sound have strong and consistent offshore wind (Manwell & Rogers, 2002).

It is becoming harder for sailors to navigate as changes to environmental conditions make less predictable wind, seas, and routes. Climate change has already influenced traditional sailing routes and sailors are more likely to have to adjust their route (Singleton et al., 2023). To have successful route planning you need to understand the environmental conditions that influence a trip. As wind patterns become more variable, the ability to anticipate changing conditions will also likely become harder.

Although current research shows the changing wind patterns and coastal environments, a lot of the existing literature focuses on either global climate trends or broader regional processes. Very few studies specifically look at how these environmental changes could influence wind conditions around Nantucket from the perspective of sailing and safe navigation.

The purpose of this literature review is to look at how climate change and other environmental changes may influence coastal wind and wave patterns near Nantucket and the effects of these changes on sailing conditions. To find this, peer-reviewed journal articles, emphirical papers, and papers that spoke to my research question were found through google scholar using the search terms "climate change" AND "wind speed" AND “sailing” AND "Nantucket".

Methods

This literature review was conducted to look at existing research on how changes in wind speed, wave heights, and coastal storms could affect sailing near Nantucket, Massachusetts. I

used Google Scholar to find relevant literature using search terms "climate change" AND "wind speed" AND "Nantucket" AND "wave height" with a five year time restriction from 2021 to 2025. The search resulted in 44 readings. My inclusion criteria were (1) studies had to be peer-reviewed, (2) must be empirical papers, and (3) discuss changes to climate and coastal regions that relate to my research question. After judging the literature with these requirements, I selected nine papers to include in the review.

Results

All 9 of the peer-reviewed empirical papers looked at environmental processes that influence wind, waves, or coastal storms in the Northwest Atlantic or close coastal regions. Most studies focused on the northeastern United States, including Massachusetts coastal waters, the Gulf of Maine, and the Northwest Atlantic, while several looked at offshore wind areas near south of Massachusetts and Rhode Island, all of which are relevant to Nantucket.

Additionally, many of the studies used numerical modeling instead of just observational data, such as atmosphere-ocean-wave modeling systems. These are used to simulate real world interactions between waves, wind, tides, and sea-level rise. Although the methodologies were different across the papers, they all examined similar variables like wind speed, wave height, storm surge, currents, surface temperature, and sea level rise.

Changes in wind patterns

The literature shows how changes in wind patterns affected ocean circulation, atmospheric stability, and wave development. Two studies looked at anthropogenic changes with offshore wind farms, while others found that during storms, wind interacts with waves, tides, and currents to affect the marine environment (Seo et al. 2025; Sauvage et al. 2025; Lin et al. 2021; Lin and Sheng 2023). Seo et al. (2025) discovered that large-scale offshore wind farms reduced hub-height wind speeds (wind speed measured at the height of the turbine) by approximately 2–3 m s^{-1} (20–30%) and 10-m wind speeds by 0.25–0.50 m s^{-1} (5–10%) downstream of the area (Seo et al. 2025). Meaning, the offshore wind turbines slow down the wind after it passes through the wind farm, creating a slower ocean breeze near the wind farm. Further, the reduction in wind speed/stress (the force that the wind exerts on the oceans surface) decreased upper-ocean mixing, which led to a 0.3 -- 0.4°C increase in sea-surface temperature (Seo et al. 2025). These changes combined also increased the variability of atmospheric conditions by about 5%, a pretty large percentage for such small changes in wind force and speed (Seo et al. 2025).

Similar patterns were found in wind speed reduction by Sauvage et al. (2025), who looked at how offshore wind farm wakes changed wave development along the northeastern United States coast. Their simulations showed that surface wind speeds decreased by about 10% downstream and within the area, reducing wave-supported momentum flux (the force the wind puts on the ocean to help create waves) by approximately 30% (Sauvage et al. 2025). As wind energy and force that was going on the oceans surface declined, wind-generated waves also decreased. This caused longer-period swells to become increasingly larger as they made up a

larger share of total waves. (Sauvage et al. 2025). Together, these studies demonstrate how even small changes to coastal wind speed can have significant impacts on both atmospheric and marine processes (Seo et al. 2025; Sauvage et al. 2025).

Although the remaining studies did not directly look at long-term changes in wind speed, they proved that wind can't be considered separate from other coastal processes. Lin et al. (2021) showed that interactions between hurricane winds, ocean currents, and surface waves changed wave heights by more than 10% and increased storm surge by as much as 19% (Lin et al. 2021). Likewise, Lin and Sheng (2023) found that when waves and currents interact, the speed/direction of the water at the surface changed by up to **0.55 m s^{-1}**, making waves about 20% smaller because of the currents(Lin and Sheng 2023). Together, these sources show how changing wind conditions have effects beyond the atmosphere and can actually influence marine processes through interactions with waves, tides, and ocean circulation (Lin et al. 2021; Lin and Sheng 2023).

Changes in Wave and Sea Conditions

The literature also found that environmental changes influence wave height, wave energy, storm surge, and overall sea-state conditions throughout the Northwest Atlantic. Although the impact of these conditions changed depending on storm intensity and atmospheric conditions, all of the selected papers found that waves are influenced by wind, tides, ocean currents, and sea-level rise, not just wind force alone (Chen et al. 2021; Lin et al. 2021; Lin and Sheng 2023; Hallowell et al. 2021).

A few studies found that wave and current interactions during storms can significantly change wave properties, like height, speed, and direction. Lin et al. (2021) compared models that included the interaction between waves and ocean currents with models that measured waves separately. When they included wave–current interaction in the measurement they predicted maximum wave height was more than 10% lower on the path of hurricanes. So, ignoring ocean currents can cause a model to overestimate how high the waves get during hurricanes. Similarly, Lin and Sheng (2023) found that in places where ocean currents significantly change in a short distance, wave–current interactions caused wave heights to decrease by about 2.1 m (about 20%), and at the same time increase surface ocean currents by about 0.55 m s^{-1}. Both studies found that accurate predictions of ocean conditions need models that look at wave and current interactions together instead of separately.

It was also found that sea level rise can intensify wave-related coastal hazards. Chen et al. (2021) found that projected sea-level rise increased wave heights by 0.1–1.0 m in Boston Harbor and 0.2–1.2 m in Scituate Harbor by using coupled hydrodynamic simulations of the 1978 Massachusetts Blizzard (Chen et al. 2021). The authors even claimed that, "wave runup-produced overtopping has the same order of magnitude as surge-induced flooding" (Chen et al. 2021).

Additionally, Hallowell et al. (2021) found that the water southeast of Cape Cod and around Nantucket are vulnerable areas to hurricane-generated breaking waves because of

shallow bathymetry. Their model showed that 29% of Atlantic offshore wind development areas exceeded breaking-wave thresholds (Hallowell et al. 2021).

Implications for sailing

Sailing relies on the interaction between wind and water, so changes in wind speed alone do not determine sailing conditions. The reviewed literature showed that winds interact with waves, currents, tides, water depth, and atmospheric conditions, and changes to one thing can influence several others (Lin et al. 2021; Lin and Sheng 2023; Seo et al. 2025). This is particularly relevant near Nantucket because Hallowell et al. (2021) found that waters southeast of Cape Cod and beyond Nantucket are areas of elevated breaking-wave hazard during hurricanes. These conditions are important to sailors because safe navigation and route planning depend on both wind and sea conditions (Singleton, Pickhardt, and Young 2023).

Changes in coastal winds may also affect the conditions sailors have after extreme storms. Seo et al. (2025) used models and found wave reductions of 5–10% in 10-m wind speeds near offshore wind developments, while Sauvage et al. (2025) found about 10% reductions in near-surface winds, and a 3% decrease in significant wave height and a shift toward longer-period swell. These conditions and relationships become more complex during storms. Lin et al. (2021) found that wave-current interactions changed maximum wave heights by more than 10% along storm tracks and Lin and Sheng (2023) found wave heights changed up to 2.1 m (approximately 20%) during Hurricane Earl.

Longer-term environmental changes can have further effects on these hazards. Chen et al. (2021) found that sea-level rise increased wave heights during simulations of a major Massachusetts storm, and Hallowell et al. (2021) showed that shallow waters near Nantucket can increase susceptibility to breaking waves. However, the reviewed studies primarily investigate physical environmental processes rather than sailing itself (Chen et al. 2021; Hallowell et al. 2021). Although sailing depends on environmental conditions like wave height, current, and storms, the literature supports the idea that environmental changes are altering conditions that affect sailing. However, it doesn't say exactly how these changes are affecting sailing safety or performance near Nantucket.

Limitations

One limitation of this review was that many of the reviewed studies used models to simulate waves, currents, storms, or environmental changes, including Chen et al. (2021), Lin et al. (2021), Lin and Sheng (2023), Seo et al. (2025), and Hughes et al. (2025). Hughes et al. (2025), and relied on numerical data instead of long term observations. Further, the reviewed studies provide evidence about changing winds, waves, currents, sea-level rise, and storm hazards, but don't directly measure how these changes affect sailing safety, performance, or navigation near Nantucket (Chen et al. 2021; Lin et al. 2021; Hallowell et al. 2021; Lin and Sheng 2023). So, the evidence shows that the marine environment changes affect importable variables that are necessary for sailing, but more Nantucket-specific research is needed.

Conclusion

Sailing depends on sailors' ability to understand and read waves, winds, and changing sea conditions. However, as the environment continues to change, understanding the relationships and how these variables work together will become increasingly difficult for sailors along the Northeast U.S. coast. This review found that changes in wind can affect the waves and ocean–atmosphere interactions. Further, the review found that sea-level rise, storms, currents, and coastal bathymetry can create more complex and potentially dangerous sea conditions (Chen et al. 2021; Lin et al. 2021; Hallowell et al. 2021; Lin and Sheng 2023; Seo et al. 2025). For Nantucket specifically, these findings have implications for recreational sailors, competitive sailors, and the coastguards who have to plan and manage marine routes and activities. The reviewed literature also showed that future conditions around Nantucket will be influenced by both climate and human related environmental changes. For example, offshore wind development can modify wind speeds, wave heights, and ocean–atmosphere interactions and the expansion of offshore wind farms should be considered with climatic changes when thinking about future marine conditions (Seo et al. 2025; Sauvage et al. 2025).

Works Cited

Anderson, Byron D. The Physics of Sailing Explained. Sheridan House, 2003. https://books.google.com/books?hl=en&lr=&id=kFjzekyPKe4C&oi=fnd&pg=PA1&dq=sailing+and+wind+patterns,+why+so+important%3F+AND+%22patterns%22+AND+%22recreation&ots=qm_0vTSEyP&sig=w-U36QJrzv70n0hhJO_oN3Fp4QA#v=onepage&q&f=false.

Archer, C. L., et al. "On the Predominance of Unstable Atmospheric Conditions in the Marine Boundary Layer Offshore of the U.S. Northeastern Coast." Journal of Geophysical Research: Atmospheres, vol. 121, 2016, pp. 8869–8885. https://doi.org/10.1002/2016JD024896

Brown, Wendell S. "Wind-Forced Pressure Response of the Gulf of Maine." Journal of Geophysical Research: Oceans, vol. 103, no. C13, 1998, pp. 30661–30678. https://doi.org/10.1029/98JC00813

Ding, Li, and Thalis P. V. Zis. "Ship Weather Routing: A Taxonomy and Survey." Ocean Engineering, 2022. https://www.sciencedirect.com/science/article/abs/pii/S0029801820306879. Accessed 27 July 2026.

Manwell, J. F., and A. L. Rogers. "An Offshore Wind Resource Assessment Study for New England." Renewable Energy, Oct. 2002. https://www.sciencedirect.com/science/article/abs/pii/S0960148101001835 Accessed 27 July 2026.

Maine Climate Council, Scientific and Technical Subcommittee. Scientific Assessment of Climate Change and Its Effects in Maine: 2024 Update. Maine Climate Council, 2024.

Schotte, Margaret E. Sailing School: Navigating Science and Skill, 1550–1800. Johns Hopkins University Press, 2019.

Singleton, Frank, Fred Pickhardt, and Jared Young. Climate Change and Ocean Cruising. Ocean Cruising Club, May 2023. https://liveicomcontent.azureedge.net/occfiles/Uploadedpdfs/20230703_REVIEW%20v.2.5%20Lowres%202.pdf. Accessed 20 July 2026.

Wang, O., et al. "Local and Remote Forcing of Interannual Sea-Level Variability at Nantucket Island." Journal of Geophysical Research: Oceans, vol. 127, 2022, article e2021JC018275. https://doi.org/10.1029/2021JC018275

Wu, Jian, et al. "Changes in Terrestrial Near-Surface Wind Speed and Their Possible Causes: An Overview." Climate Dynamics, 9 Nov. 2017. https://doi.org/10.1007/s00382-017-3997-y. Accessed 20 July 2026.

Zheng, Chong-wei, and Xue-hong Li. "Global Trends in Oceanic Wind Speed, Wind-Sea, Swell, and Mixed Wave Heights." Applied Energy, 1 Sept. 2022. https://www.sciencedirect.com/science/article/abs/pii/S0306261922006778#preview-section-cited-by. Accessed 20 July 2026.

How Does the Rotation Curve of a Spiral Galaxy Deviate from Keplerian Predictions, and What Does This Imply About the Distribution of Dark Matter?

By Agastya Nambiar

Abstract

The rotation curves of spiral galaxies present difficult conundrums to modern astrophysics. According to simple Newtonian mechanics, rotation velocities should decline as one moves farther away from the centre of the galaxy, giving a Keplerian falloff. In the case of spiral galaxies this is not the case. For decades' scientists have utilized both spectroscopic and radio observations to show that beyond the stellar disk, the rotation curves of galaxies remain approximately flat or rising, far exceeding the visible extent of galaxies. This paper addresses the theoretical aspects of the prediction of Keplerian dynamics and the methods used to measure the rotation curves of galaxies and the scale of the discrepancy between the predictions and the observations. Compelling evidence of dynamically dominant mass exists that is not revealed in the electromagnetic spectrum. This unseen mass is known as dark matter. Of the many models of dark matter halos, the isothermal sphere and the Navarro Frenk White profile are evaluated against the dark matter distribution in spiral, dwarf and low-surface-brightness galaxies. Alternative explanations, particularly Modified Newtonian Dynamics (MOND), are also critically examined. The paper concludes that dark matter halo models successfully explain the broad features of observed rotation curves, although important questions remain regarding the detailed structure of dark matter halos, especially in their central regions.

Keywords Spiral galaxies, Rotation curves, Keplerian dynamics, Dark matter distribution, Dark matter halos, Galactic dynamics, Navarro Frenk White profile, Modified Newtonian Dynamics (MOND).

1. Introduction

For close to one hundred years, astronomers have endeavoured to answer whether all the gravitational mass of a galaxy is represented by its visible matter. One of the first mass discrepancy cases outside the galaxy was advocated by Fritz Zwicky, who discovered that the total mass of the Coma Cluster, as determined by the velocity dispersion of the galaxies, was several hundred times more than the mass of the visible galaxies contained within the Cluster (Andernach & Zwicky, 2017; Zwicky et al., 2021). Although Zwicky's result was largely ignored, the high-quality rotation-curve data for individual spiral galaxies that began to emerge in the early 1970s put the astrophysics case for "missing mass" front and centre.

Zwicky supported his argument using the virial theorem, which states that for a system of particles obeying the fundamental law of gravity, in equilibrium, the relation ($2T + U = 0$) holds, where (T) and (U) denote the total kinetic and gravitational potential energies, respectively. By approximately determining these energies from the velocity dispersion and characteristic radius of a cluster, the mass of the cluster can be estimated using the virial theorem $M \approx rv^2/G$), a

quantity that is independent of the system's luminosity. When Zwicky applied this method to the Coma Cluster, he found its total mass to be nearly two orders of magnitude greater than the mass inferred from its visible galaxies alone (Zwicky, 1933). In contrast to the estimation of masses from emitted radiation, the virial method indicates the gravitating mass through the dynamical behavior of the system. While theoretically distinct from the galaxy rotation-curve technique discussed in this paper, the two methods provide complementary evidence for the presence of unseen matter.

In a gravitational system that obeys Newtonian dynamics, a test particle in a circular orbit at a radius r, about a mass M(r), will orbit with a velocity given by the equation,

$$v(r) = \sqrt{\frac{GM(r)}{r}}$$

For a galaxy, where the large portion of the luminous mass is contained within a finite disk and a central bulge, this relation would demand that rotational velocities would rise to the edge of the disk and would then fall away as $v \propto r^{-1/2}$, in the Keplerian regime. This result would be predicted by an analogy of the galaxy to a planetary system, where the majority of the mass would be situated at the centre (McGaugh et al., 2016).

By the mid-1980s, the rotational patterns of spiral galaxies contradicted what was expected of them, and modern surveys fully supported this. The velocity of orbiting stars and gas was found to be constant or steadily increasing, even at large distances from the centre of the galaxy. While some of that matter belonged to visible stars and gas and was luminous, much of that mass much of the mass required to explain the observed velocities was still invisible (Famaey & McGaugh, 2012). Because of this, galaxies needed to be surrounded by halos of dark matter to provide additional gravitational influence to sustain those stars' and gas clouds' orbits. Most of the non-baryonic and non-luminous matter found in the halos around galaxies is dark matter (Young, 2017).

2. Theoretical predictions: keplerian dynamics and galactic rotation

2.1 The Keplerian Velocity Equation

For a point mass or a mass distribution with a spherical symmetry, Newtonian gravitational mechanics states that circular orbital speed v at radius r would satisfy:

$$v(r) = \sqrt{\frac{GM(r)}{r}}$$

where M(r) is the mass enclosed at radius r. This expression can be obtained by equating gravitational force and centripetal force. When r is much greater than the scale of mass distribution, M(r) approaches a constant value M total, the speed of the orbiting velocity should decrease with distance i.e. $v \propto r^{-1/2}$, the so-called Keplerian rotation. The outer planets of the solar system exemplify the inverse square root. The Sun contains more than 99.8% of the mass of the solar system providing the outer planets' Keplerian rotation (Binney & Tremaine, 2008).

2.2 Application to Spiral Galaxies and the Expected Velocity Profile

Applying this framework to a spiral galaxy needs an understanding of the distribution of stellar mass. The surface brightness of a spiral galaxy disk is described well by:

$$I(R) = I_0 e^{-R/h}$$

where I_0 is the central surface brightness and h is the disk scale length. Freeman (1970) studied 36 spiral and S0 galaxies and found that the central surface brightness I_0 was almost constant in galaxies of different luminosities and sizes (Freeman, 1970). This result is referred to as Freeman's Law. If the disk has a purely exponential shape and surface brightness with an associated central disk surface brightness and scale length, then the peak of the rotation curve is located at about $R \approx 2.2h$ and the curve then declines. At large enough radii, the mass of the disk M_disk becomes approximately constant and then $v(r) \propto r^{-1/2}$ is expected.

By including a stellar bulge, modeled as a concentrated mass at the center, a rapidly rising inner rotation curve is produced that later declines steeply. Neither the disk nor the bulge contribute to the observed flat rotation at large radii. The critical prediction is that a galaxy formed from the entirety of its stellar and gaseous baryonic matter should show a Keplerian velocity decline beyond the disk edge. The decline that is expected can be determined if the stellar population's mass-to-light ratio is known. These predictions show that at large radii the observed rotation velocities exceed those that are expected from the visible mass alone (Alexeev, 2017; Famaey & McGaugh, 2012).

3. Structure and luminous mass distribution of spiral galaxies

At the very least, spiral galaxies contain three baryonic components. These components are: star and interstellar gas disks that are supported by the galaxy's rotation; bulges that contain older stars and are centrally concentrated; and diffused ionized gas that resides in the outer disk (Martinsson et al., 2013). All of these components have a unique absorption profile, a unique kinematic structure, and a unique mass-to-light (M/L) ratio. To determine the true nature and extent of any mass deficit, one must first accurately represent all the luminous mass present.

In the case of late-type spiral galaxies, the stellar disk is the component that contains the majority of baryonic mass. Following Freeman (1970), the surface density of the stellar disk is considered to have an exponential profile with a scale length (h) that can vary from 1 to 5 kpc for galaxies of different morphological types (Sofue & Rubin, 2001). The stellar disk provides a rotation curve that is characterized by a rapid rise in the inner regions of the disk and a gradual decline in the outer regions of the disk beyond approximately 2h. The curve is sensitive to the mass-to-light (M/L) ratio, which is a significant source of uncertainty in mass modelling. With old stellar populations, this ratio has a range of ($M_\odot/L_\odot$) in the R band of about 0.5 to 5 and is dependent on the initial mass function and the stellar population's star-formation history (Bell & De Jong, 2001).

The deepening of our understanding of neutral hydrogen through its 21-cm hyperfine emission line shows that the hydrogen disk in a galaxy extends beyond its stellar disk as much as 3× the optical radius. The gas disk is less massive than the stellar disk. The gas disk is observed more easily than the stellar disk as the gas does not require a conversion of luminosity to stellar

mass (which is an uncertain process). In late-type spirals, the gas disk is 5–20% of the total baryonic mass and in the inner disk its contribution to dynamics is very small compared to that of the stellar disk (Sofue & Rubin, 2001). In the outer disk however, it starts to become more important.

Where it is present, the central bulge raises the inner rotation curve steeply. In Sa and Sb spirals it dominates the mass within small radii. In Sc and Sd spirals it becomes less important. In order to accurately describe the inner rotation curve of a galaxy one must account for the mass of the bulge. However, since it does not affect the outer rotation curve, it does not change the conclusion that a substantial amount of dark matter is required in the outer regions.

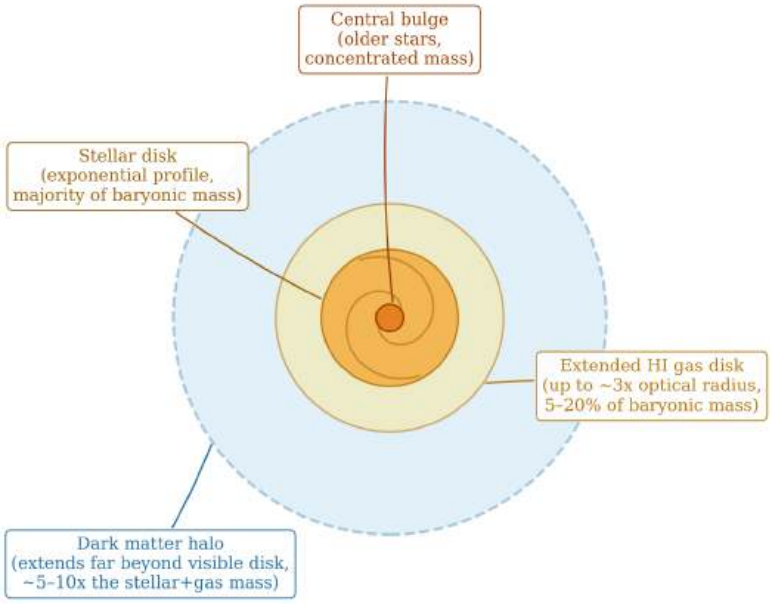


Figure 1 Galaxy mass components schematic diagram

4. Observational Methods for Measuring Rotation Curves

4.1 Optical Spectroscopy and H-alpha Emission

Rotation curves were traditionally built using optical spectroscopies of ionised gas emission lines, in particular, Hα at 656.3 nm. HII regions, which are concentrated in the spiral arms, can be used to find the velocity of gas. The Doppler shift of Hα emissions in relation to the systematic velocity of the galaxy gives the line of sight velocity. This projected velocity can be corrected using the inclination angle of the galaxy to give the circular velocity at that radius. This method was first used by Rubin & Ford (1970) in their spectroscopic study of the Andromeda galaxy, M31 (Rubin & Ford, 1970). They extended their rotation curves to 110 arcminutes and reported a rotational velocity of 200–250 km s^{-1}. They also reported an absence of a Keplerian decline for the outer disk regions.

Optical spectroscopy is limited as a technique as HII regions become sparse in the outer disk of the galaxy. This technique only covers the luminous sections of the galaxies. Beyond that optical disk, ionised gas regions are too sparse to be used as kinematic tracers.

4.2 Neutral Hydrogen 21-cm Radio Observations

Neutral hydrogen's (HI) 21 cm spin-flip frequency transition serves as a more useful kinematic tracer when HI gas extends well beyond the stellar disk. HI gas can extend beyond the stellar disk by a radius up to double the optical radius. In addition, HI gas enables the measurement of rotation curves at larger galactocentric radii because the Doppler-shifted 21 cm line can easily be detected using radio telescopes and the aerial coverage of HI gas is extensive. Roberts & Whitehurst (1975) were able to extend the M31 rotation curve out to ~30 kpc using HI 21 cm observations and found that the outer rotation of HI remained flat and unchanging (Roberts & Whitehurst, 1975). Bosma (1981) performed an HI survey on 25 spiral galaxies and concluded that flat rotation curves were a ubiquitous characteristic among spiral galaxies of different morphological types (Bosma, 1981).

The use of aperture synthesis interferometry in modern 21 cm observations enables the detailed study of HI gas and the gas kinematics in nearby galaxies using radio telescope arrays with even greater angular resolution. The HI Nearby Galaxy Survey (Walter et al., 2008) was able to produce unprecedented data for the analysis of galactic rotation curves by 21 cm mapping 34 nearby galaxies with even greater resolution and sensitivity than prior studies. The SPARC database (Lelli et al., 2016)also provided accurate HI and Hα rotation curves for 175 disk galaxies, and statistical studies in that database spanned six orders of magnitude in stellar masses and addressed the interactions between dark matter and luminous matter.

In HI interferometry, an observational issue is beam smearing. Finite angular resolution causes an averaging of velocities at small radii which results in a systematic underestimation of the inner rotation curve's slope. That's why this phenomenon mimics the presence of a constant-density (cored) dark matter halo. For dwarf galaxies especially at the moderate angular resolutions typical in HI interferometry, this effect y impacts the measurement by a lot and dark matter halo profiles cannot be estimated in a reliable manner until this is properly accounted for (De Blok, 2010).

5. Observed rotation curves: the systematic violation of keplerian predictions

5.1 Historical Discovery and Systematic Evidence

The failure of the Keplerian prediction stemmed from the gradual development of consistent observations that grew in accuracy over about 10 years, as opposed to a single major event. According to Rubin and Ford (1970), the rotation curve for M31 was flat and showed no evidence for the anticipated decline, and their finding was of little concern for most. Roberts and Whitehurst (1975) advanced this finding to even greater distances in M31, measuring rotation

velocities of about 225 km s^{-1} at distances of about 30 kpc, which is much further out than the stellar disk.

The case for flat rotation curves as a general phenomenon was made by Rubin, Thonnard, and Ford in their studies of late-type spirals. In their 1980 paper, which compared 21 of Sc galaxies spanning across a luminosity range ~50, all of the galaxies surveyed had roughly flat rotation curves that were located between two disk scale lengths and the largest measured radii (Rubin et al., 1980). The galaxies' rotation was flat and the velocities varied between the least luminous at ~90 km s^{-1} and the most luminous at ~300 km s^{-1}, showing that a flat rotation curve is a universal characteristic and not a feature of one specific mass scale. With his 21-cm survey, Bosma (1981) found that flat HI rotation curves were found in all spiral galaxies, and that, at 2-3 optical radii, they showed no evidence of decline.

A quantitatively simple example illustrates the difference between the predicted Keplerian rotation and what is observed. For a galaxy with a total stellar mass ~10^{11} $M_\odot$ and a distribution of stellar mass concentrated within ~10 kpc, Keplerian dynamics yields a prediction of $v \approx 65\ (M/10^{11}\ M_\odot)^{1/2}\ (r/30\ \text{kpc})^{-1/2}$ km s^{-1}, indicating a rotational speed of roughly 65–100 km s^{-1} at a distance of 30 kpc. In contrast, the same distance in similar galaxies yields $v \approx$ 200–250 km s^{-1}, which suggests that the mass contained within a sphere of 30 kpc is approximately 5–10 times the total stellar mass. The mass in the visible components of the galaxy is insufficient to support the observed rotations (Jiao et al., 2023; Sofue, 2013).

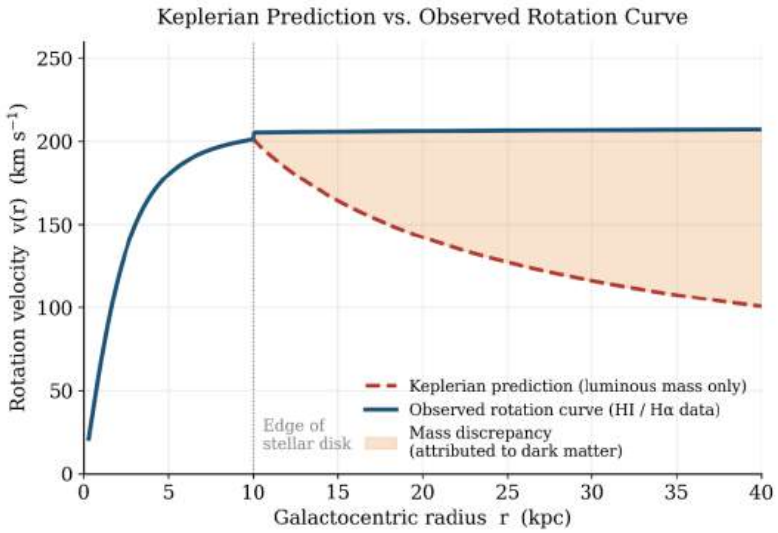


Figure 2 Schematic comparison of Keplerian and observed rotation curves

5.2 The Universal Rotation Curve

According to research by Persic, Salucci, and Stel (1996), analyzing almost 1100 spiral galaxies demonstrated that low luminosity spiral galaxies had a rotation curves which slowly increase and do not yet meet a flat asymptotic value at the largest observed radii (Persic et al., 1996). High-luminosity spiral galaxies have a steep increase in rotation due to the stellar bulge

with a declining curve that is not as steep as the Keplerian decline. Intermediate luminosity galaxies have a flat rotation curve. These findings suggest that the ratio of dark to luminous mass is a function of spiral galaxy mass which acts as a constraint on dark matter halo models.

5.3 The Tully–Fisher Relation as Corroborating Evidence

Empirically relating the total luminosity of a spiral galaxy to its maximum rotation velocity, the Tully-Fisher relation supports the theory of the existence of some regulated unseen mass that influences galactic dynamics. If rotation velocities were proportional to the stellar mass present, then the Tully-Fisher relation would be expected to be a scattered relation, due to the mass-to-light ratios of galaxies' stellar components differing to such an extent. The strong correlation of the Tully-Fisher relation is evidence that total gravitating mass, including dark matter, does not deviate significantly from the extent of a galaxy's luminous mass and is closely related. This constrains the formation of galactic halos and indicates to a noteworthy degree, a pattern for the ratio of baryonic and dark matter (Sofue & Rubin, 2001).

6. Dark matter halos: density profiles and mass models

6.1 The Isothermal Sphere

The singular isothermal sphere is the simplest model that can create a flat rotation curve. Over here the dark matter density falls off as $\rho(r) \propto r^{-2}$. A better model is the non-singular pseudo-isothermal (PI) profile which is expressed as:

$$\rho(r) = \frac{\rho^0}{1+r^2/r_c^2}$$

where ρ_0 is the density at the center and r_c is the core radius (Begeman et al., 1991). This profile has a central core that is flat and has a density that is approximately constant. At large distances, this profile has a density that falls off as $\rho \propto r^{-2}$ and leads to a flat rotation curve. The velocity that is associated with this curve is:

$$v^2(r) = 4\pi G\rho_0 r_c^2 [1 - \frac{r_c}{r}\arctan\ arctan\ (\frac{r}{r_c})$$

An asymptotic velocity $v_\infty = \sqrt{4\pi G\rho_0 r_c^2}$ is achieved at large distances. The PI model is a better fit for rotation curves and has the advantage of being simple to parameterize, as it requires only two parameters which can be adjusted based on the data.

An illustration of the method can be found in the way van Albada et al. (1985) treat NGC 3198 in the canonical rotation curve decomposition. From photometric data, the stellar component of the rotation curve is calculated, and the gaseous component is obtained from 21-cm maps . The total rotation curve is then fit with a pseudo-isothermal dark matter halo, where the parameters ρ_0 and r_c are adjusted. Van Albada et al. showed that a dark halo exceeding 80% of the total mass at the outermost observed radius was required to model the HI rotation curve of NGC 3198, which is flat at $v \approx 150$ km s^{-1} out to ~30 kpc. Their work was one of the

first to build a strong case that dark matter must dominate the large-scale gravitational potential of a spiral galaxy at large radii (Van Albada et al., 1985).

6.2 The Navarro-Frenk-White Profile

An important theoretical development was made through cosmological N-body simulations exploring structure formation generated with Cold Dark Matter (CDM). These simulations predicted that dark matter halos would develop a characteristic density profile that was cuspy and that this density profile would differ from the cored pseudo-isothermal profile. Navarro, Frenk & White (1996, 1997) demonstrated that in numerous simulations of CDM across many orders of magnitude in mass, dark matter halos could be modeled using the density profile given by:

$$\rho(r) = \frac{\rho_{crit}\delta_c}{\left(\frac{r}{r_s}\right)\left(1+\frac{r}{r_s}\right)^2}$$

where r_s is scale radius, ρ_{crit} is the critical density of the universe, and δ_c is a dimensionless overdensity (Navarro et al., 1996, 1997). This profile demonstrates a steep rise with a slope of $\rho \propto r^{-1}$ and at large distances from the center, this profile falls as $\rho \propto r^{-3}$. For NFW halos, the concentration parameter is,

$$c = \frac{r_{200}}{r_s}$$

r_{200} is the virial radius that represents the mass assembly history of the halo. More concentrated halos are those that formed earlier in cosmic history (Navarro et al., 1997). The NFW profile produces a rotation curve that rises to a broad peak and then declines, which agrees with observed rotation curves that are extended and flat for large spiral galaxies, though this profile is less appropriate for the rising rotation curves that are observed in small dwarf galaxies.

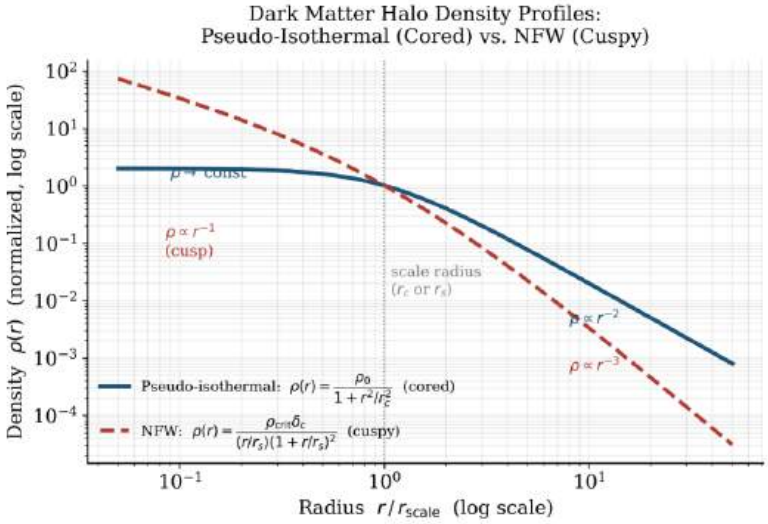


Figure 3 Comparison of pseudo-isothermal (cored) and NFW (cuspy) dark matter density profiles

6.3 The Burkert Profile

A widely accepted core-halo model is the Burkert profile. This profile was developed empirically to match the observed rotation curves of low-surface-brightness and dwarf galaxies (Burkert, 1995). It can be expressed as:

$$\rho(r) = \frac{\rho_0 r_b^3}{(r+r_b)(r^2+r_b^2)}$$

Here ρ_0 denotes the central density and r_b is the core radius. Similar to the pseudo-isothermal profile, it predicts a constant-density core that is equal to ρ_0 as r □ 0. However, at large radii the density decreases as $\rho(r) \propto r^{-3}$ as opposed to $\rho(r) \propto r^{-2}$ decline in density as predicted by the pseudo-isothermal sphere. As a result, the Burkert profile assumes an important place in the core-cusp debate, as it conjoins an observationally justified constant density core with the large radius density behavior predicted by cold dark matter scenarios (Salucci & Burkert, 2000). Analysis of high resolution rotation curves for a number of spiral galaxies, including NGC 3198, have generally shown that cored models, for example, the Burkert and pseudo-isothermal profiles, do provide better fits to the inner regions as compared to the cuspy NFW profile, thus strengthening the evidence for the core-cusp discrepancy.

6.4 The Core Cusp Debate

The difference of the cored PI profile (preferred by observations) versus the cuspy NFW profile (attributed to CDM simulations) has led to the so-called "core-cusp problem". High resolution rotation curves of low luminosity and dwarf irregular galaxies where the dark matter contribution is less contaminated by stellar mass, show uncoupled dark matter and stellar mass and are better fit with core models rather than NFW cusps. Moore et al. (1999) argued CDM simulations predict an even steeper cusp than NFW of the order of $\rho \propto r^{-1.5}$, increasing the discrepancy (Moore et al., 1999). De Blok (2010) stated that observational data generally favour inner density profiles that are shallower than the NFW cusp, even though the precise slope of the inner profile remains uncertain due to beam smearing, non-circular motions and projection effects. Recent simulations that use baryonic physics (and supernova feedback and gas outflows) show that repetitive cycles of gas inflow and outflow can transform cusps into cores. This may be a solution to the CDM problem (de Blok, 2010).

A high-resolution study of NGC 3198 performed by Karukes, Salucci, and Gentile (2015) using data from the HI kinematics that extend to about 48 kpc, provides a good example of this core-cusp discord. In their analysis, it was found that both the cored Burkert halo and the cuspy NFW halo yielded good values of reduced chi-square ($\chi^2_r \approx 0.9$ and $\chi^2_r \approx 0.8$ respectively). The results show that, even though the shape of the rotation curve was fitted conventionally, a classic example of rotation curve fitting degeneracy occurred in which a single, simple and flat rotation curve was not sufficient to determine conclusively the presence of a core or a cusp in the NFW halo model. To overcome this problem, the authors employed an independent method that directly reconstructed the dark matter density instead of a global halo fit, using the local shape of

the rotation curve. This analysis yielded a density profile that was inconsistent with the NFW model and was more compatible with a halo of constant density. Additionally, they found that, for a conventional fit to the rotation curve, the reference (standard) concentration-mass relation for the halo of comparable mass yielded a lower value of concentration ($c \approx 6.7$) for the best-fitting NFW halo, when compared to the results of other studies. Together, these results provided evidence beyond the goodness-of-fit statistics alone, supporting the presence of a core for NGC 3198 (Karukes, Salucci, & Gentile, 2015).

7. Evidence from diverse galaxy populations

7.1 Low-Surface-Brightness Galaxies

Low surface brightness (LSB) galaxies are disk galaxies that have low stellar surface density. They are approximately an order of magnitude lower than that of normal spiral galaxies at similar luminosities. Because LSB galaxies have a low stellar mass, they have a high dark matter to stellar mass ratio. The dark matter contribution is expected to dominate the total rotation curve at nearly all radii, which diminishes uncertainties tied to the stellar mass to light ratio. De Blok & McGaugh (1997) studied the rotation curves of LSB galaxies and they found that LSB galaxies have rotation curves that rise slowly and are well fit with dark matter halos that are low in density (De Blok & McGaugh, 1997). These fits are consistent with dark matter halos that are cored or of low concentration. The mass models suggested that LSB galaxies contain greater than 90% dark matter in the mass within the optical disk. This also supported that LSB galaxies are some of the most dark matter dominated galaxies in the universe. The rotation curves of LSB galaxies cannot be fit with a purely baryonic mass model. LSB galaxies require a predominantly dark matter halo that is consistent with, but not exclusively, CDM of dark matter halo models (Di Paolo et al., 2018).

7.2 Dwarf Irregular Galaxies

Dwarf irregular galaxies remain some of the best candidates for studying the structural formation of dark matter halos. Their dominance of dark matter extends to all radii, and velocity dispersion as well as velocity rotation are similar, indicating that pressure support should be considered in their mass estimation. In addition, their simple shapes leave little room for error from non-circular motions. In the THINGS HI study, Oh et al. (2011) used high resolution rotation data to generate rotation curves for seven nearby dwarf irregular galaxies after accounting for asymmetric drift, which is the effect of kinematic pressure support resulting in a lesser observed rotation in comparison to the circular velocity (Oh et al., 2011). They argued five of the seven galaxies showed a preference for core-halo models over the NFW profile, supporting core-cusp tension. They noted, however, that beam smearing, as well as the treatment of velocity dispersion, had a profound effect on the determined inner slope, and that the evidence for cusp haloes was not as conclusive as a straightforward comparison of profile fits would indicate.

7.3 The Milky Way

Constructing the rotation curve of our own Galaxy is more difficult than for external galaxies because the observer is within the disk and measurements of distance are required for every kinematic tracer. However, many independent studies utilizing varied tracers such as HII regions, molecular masers, and red giant stars, have yielded a rotation curve of the Milky Way which is roughly the same as external spiral galaxies. Sofue (2012) synthesized a rotation curve of the Milky Way extending to about 200 kpc with a rotation speed of about 200 km s^{-1} in the outer disk, suggesting a dark matter halo of a total virial mass of about 2×10^{12} $M_{\odot}$ (Sofue, 2012). Reid et al. (2014) estimated the rotation speed of the Galaxy at the Solar circle to be $\Theta_0 = 240 \pm 8$ km s^{-1} via the trigonometric parallaxes and proper motions of star forming regions, measured with the Very Long Baseline Array (Reid et al., 2014). This also implies a rotation curve of the Galaxy which is not of the Keplerian nature, and indicates a large amount of dark matter beyond the stellar disk.

7.4 The Radial Acceleration Relation

McGaugh, Lelli, and Schombert (2016) constructed a correlation between the observed centripetal acceleration in the galactic disk and the acceleration predicted from the distribution of the galactic baryonic mass. This correlation spanned six decades of acceleration. From the SPARC database, they compiled a sample of 153 galaxies. They asserted that there was a tight correlation between the centripetal acceleration observed, denoted as g_{obs}, and the Newtonian acceleration related to baryonic matter, denoted as g_{bar} (Lelli et al., 2016; McGaugh et al., 2016). This correlation was shown to have a scatter of only 0.11 decibels. This is known as the Radial Acceleration Relation (RAR) and states that the amount of dark matter present in a galaxy is completely dependent upon the amount of baryonic matter existing in the galaxy. This is a strong indicator of the relationship between luminous matter and dark matter, highlighting the interactions between dark matter and the baryonic matter in the process of dark matter halo formation. This Radial Acceleration Relation is a benchmark for multiple dark matter theories and for Modified Newtonian Dynamics (MOND) theory. In that sense, it is used as a benchmark to evaluate different theories (McGaugh et al., 2016).

8. Modified newtonian dynamics: an alternative interpretation

8.1 The MOND Framework

Modified Newtonian Dynamics (MOND) was introduced by Milgrom (1983) and is an alternative explanation for the flattening of rotation curves that does not rely on the presence of dark matter, instead modifying gravitational laws in instances of low acceleration (Milgrom, 1983). According to MOND, the relation of gravitational acceleration g and Newtonian acceleration gN is described by the equation:

$$\mu\left(\frac{g}{a_0}\right)g = gN$$

where $a_0 \approx 1.2 \times 10^{-10}$ ms^{-2} is constant acceleration (Milgrom, 1983; Sanders & McGaugh, 2002) and where the function μ satisfies $\mu(x) \rightarrow 1$ and $\mu(x) \rightarrow x$ in the limits of large and small x respectively. In regions of low acceleration, MOND asserts that the circular velocities in a galaxy do not vary but are given by:

$$v^4 = G M_{bar} a_0$$

where M_{bar} is the total baryonic matter in the galaxy. This also leads to the prediction of the rotation velocity for a galaxy known as the MOND Tully-Fisher relation and requires no additional free dark matter parameters.

8.2 Successes and Failures

MOND has been very successful at describing some properties of individual galaxies. For example, Begeman, Broeils & Sanders (1991) showed that MOND, with a_0 as the only free parameter, is able to fit the rotation curves of ten different spiral galaxies with comparable success to dark matter halo fits that require multiple parameters (Begeman et al., 1991). Sanders & McGaugh (2002) showed that the rotation curves of approximately one hundred galaxies including high surface brightness spiral and dwarf irregular galaxies, are well described by MOND (Sanders & McGaugh, 2002). The theory also accounts for the Tully-Fisher relation and the observed relationship between surface brightness and the shape of a rotation curve.

Though MOND does well at the galactic scale, it does not beyond that. Galaxy clusters including Bullet clusters exhibit mass to light ratios greater than 100 $M_\odot/L_\odot$. This indicates that even with MOND, modifying gravitational dynamics cannot fully account for the mass discrepancy for galaxy clusters (Sanders & McGaugh, 2002). The anisotropy of the CMB and the large-scale structure of the universe can be easily achieved in the ΛCDM dark matter theory, but require large changes in MOND-based theories. As has been noted in the relativistic extensions of MOND such as TeVeS (Bekenstein, 2004), they also have difficulties explaining some situations. Although MOND has shown remarkable success at the scale of galaxies, the community is in agreement that it lacks the ability to supplant dark matter in cosmology (Sanders & McGaugh, 2002).

9. Uncertainties, limitations, and model degeneracies

9.1 Stellar Mass-to-Light Ratios

For mass modelling of rotation curves, the largest source of systematic uncertainty is the stellar mass-to-light ratio, Y. The conversion of observed photometric luminosity to mass requires modelling of the stellar population and therefore Y is sensitive to the chemical composition, age, and the assumed initial mass function and the star formation history (Bell and de Jong 2001). The uncertainty in Y introduces uncertainty in the stellar mass contribution to the rotation curve and, therefore, in the estimate of the dark matter content. To minimize the effect of the stellar mass-to-light ratio, modern research uses optical and near-IR photometry. The near-IR light traces the stellar mass more closely. In fact, Y varies in the near-infrared by a factor of 2,

compared to the optical where it varies by an order of 10. Of course, a factor of 2 uncertainties in the stellar mass-to-light ratio still allows freedom in the halo models and contributes to the disk-halo degeneracy where a more massive stellar disk partially offsets a less massive dark halo (van Albada et al. 1985).

9.2 Inclination Angle Estimation

The conversion of observed line of sight velocity to true circular speed relies on the galaxy inclination angle *i*, and is given by the equation ,

$$v_{circ} = \frac{v_{obs}}{sin(i)}$$

For edges on galaxies where the inclination is near 90°, the correction to circular speed is small and easy to compute. In contrast for the face on galaxies (i □ 0°), the small uncertainty in i may lead to large uncertainties in v_{circ}, which may exceed several factors. Typically, inclinations are determined from the ratio of the optical or HI disk axes, assuming the disk shape is circular, which introduces a systematic bias. The THINGS survey resolved this issue by modelling the velocity field as a series of tilted rings, allowing the disk position angle and inclination to vary at each ring (Walter et al. 2008). This added more complexity to the model and resulted in parameter degeneracy.

9.3 Distance Uncertainties and Non-Circular Motions

The measurement of the distance to a galaxy places a direct constraint on the physical scale of the rotation curve, and subsequently the mass scaling. For the bulk of nearby galaxies which have been the focus of survey-based rotation curve studies, distance uncertainty is at the level of 5–15%. The mass ($M \propto v^2r/G$) also scales with radius. This means that a distance uncertainty of 10% translates to a mass uncertainty of 20% which is troublesome for modelling halo mass. While spiral and barred galaxies are known to display non-circular motions, these also cause the gas velocity field to deviate substantially from circular motions, particularly within the inner disk. If not properly addressed, these effects may cause non-circular motions which can cause local features on a rotation curve that appear to be due to large structural changes in the halo (Sofue & Rubin 2001).

9.4 Model Degeneracies

Another limitation is the disk halo degeneracy. There exist a wide variety of disk stellar mass to light ratios and halo density profiles that are still more or less compatible with a rotation curve within the bounds of observational errors. An example of this is Van Albada et al. (1985), who described maximum disk models where Y was assigned a value as high as what could be allowed dynamically. They described how maximum disk models still involved a dark matter halo dominating at larger radii, despite the fact that the disk could model the inner rotation curve completely. The opposite is true of sub-maximum disk models, where Y is assigned a value lower than what is allowed dynamically. Dark matter contributes significantly even in the inner disk. In order to distinguish between these scenarios, additional constraints other than the

rotation curve are required. An example of this could be the vertical kinematics, the bar pattern speeds or gravitational lensing. This data is lacking in most of the galaxies that have been analysed in the large rotation-coupling surveys. Lelli et al. (2016) stated that the Radial Acceleration Relation tight clustering supports that the stellar mass to light ratios are in a narrow band close to the maximum disk solutions but this is still a developing area.

10. Broader cosmological implications

There are wider implications for cosmology and ΛCDM standard cosmology when considering that spiral galaxies are well established to be embedded in dark matter halos as a result of the strong evidence provided by galaxy rotation curves (Katz et al., 2016). In the last cosmological analysis, Planck Collaboration (2020) gave a new estimate for the composition of the universe. With Λ at roughly 69.3%, dark energy gained a larger portion than 26.8% of the total composition of the universe leaving (ordinary baryonic) matter with ~4.9% and the rest for CDM (cold dark matter) (Zhang, 2024). The universe's CDM density is in nice agreement with (and independent) the dark matter fractions attributed to galaxy rotation curves (Primack, 2024). This brings early-universe cosmology and extragalactic dynamics to a remarkable agreement.

The NFW profile puts CDM structure formation on a more quantitative basis and gives a framework for predictions for halo density and concentration as a function of mass and redshift. Inner density slopes are still debated (de Blok, 2010), but large galaxy surveys and weak gravitational lensing constitute evidence that average dark matter galaxy profiles (and CDM predictions) are broadly consistent with CDM predictions in the outer regions. Furthermore, a large body of evidence supports the idea that CDM (and dark matter) is a physical reality of the universe. This evidence includes CDM and dark matter's success in describing a large-scale structure of the universe, including the cosmic web of filaments and voids, galaxy clustering and the distribution of galaxy clusters.

Currently, we understand very little about dark matter. The rotation curve depicts the gravitational effect of dark matter but cannot tell us about the nature of dark matter. There are many possibilities of what dark matter could be. They include WIMPs, sterile neutrinos, axions, and primordial black holes. Recently, the interest in WIMPs has dwindled because of the lack of signal from WIMPs. This change in focus has not affected the evidence of the gravitational impact of a non-baryonic dark matter in galaxies that rotation curves present. The legacy of using rotation curves also provides evidence of the need for more physics to be included in the Standard Model.

11. Conclusion

The rotation curves of spiral galaxies show that the gravitational mass of galaxies is much greater than the luminous baryonic content of galaxies. The Keplerian prediction shows that orbital velocities should decline with distance beyond the stellar disk, but in every spiral galaxy that has been studied, velocities remain flat at around 100-300 km/s. The discrepancy between the predicted velocities of spiral galaxies shows that outer regions of the galaxy contain more

mass than previously thought and shows how mass is distributed throughout the galaxy. In order to figure out how the rotation curves behave, mass models take into account every point on the rotation curve from the inner regions of the galaxy to the outermost regions. In every model, outer dark matter halos of the galaxy contain 5 to 10 times more mass than all the stellar and gaseous matter of the galaxy combined.

Dark matter is also needed to account for the mass of dwarf irregular galaxies and low-surface-brightness galaxies, as well as the Milky Way galaxy. The NFW cusp profile and the cored isothermal sphere can describe and predict the flat rotation curves of many galaxies and predict mass that is distributed in a certain way throughout the galaxy, but observations and theory continue to differ regarding the central density structure of dark matter halos. This discrepancy shows that there are limitations in our understanding of the baryonic content in halos.

Modified Newtonian Dynamics formulates an elegant explanation for galaxy-scale rotation curves, but it cannot explain rotation curves at the scales of clusters or the universe at large. As a phenomenological approach, Modified Newtonian Dynamics retains applicability, but dark matter remains the dominant accepted explanation.

There are many sources of uncertainty in measuring dark matter halos, such as stellar mass-to-light ratio, motion deviations from a circular orbit, beam-smeared images and halo model degeneracies. Despite these uncertainties, it is widely accepted that the majority of the mass of a galaxy exists in a dark matter halo that is both extended and gravitationally dominant. This claim is robust with respect to systematic effects and is supported by other independent studies of dark matter which include those of gravitational lensing, cluster dynamics, and the formation of large-scale structures in the cosmos. However, the distribution of dark matter, the nature of the particles that compose it, and how it came to exist as it does is one of the premier challenges currently faced by physics and cosmology.

Works Cited

Alexeev, B. V. (2017). The Dark Matter Problem. In Nonlocal Astrophysics (pp. 141–176). Elsevier. https://doi.org/10.1016/B978-0-444-64019-2.00004-1

Andernach, H., & Zwicky, F. (2017). English and Spanish Translation of Zwicky's (1933) The Redshift of Extragalactic Nebulae (Version 1). arXiv. https://doi.org/10.48550/ARXIV.1711.01693

Begeman, K. G., Broeils, A. H., & Sanders, R. H. (1991). Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics. Monthly Notices of the Royal Astronomical Society, 249(3), 523–537. https://doi.org/10.1093/mnras/249.3.523

Bekenstein, J. D. (2004). Relativistic gravitation theory for the modified Newtonian dynamics paradigm. Physical Review D, 70(8), 083509. https://doi.org/10.1103/PhysRevD.70.083509

Bell, E. F., & De Jong, R. S. (2001). Stellar Mass‐to‐Light Ratios and the Tully‐Fisher Relation. The Astrophysical Journal, 550(1), 212–229. https://doi.org/10.1086/319728

Binney, J., & Tremaine, S. (2008). Galactic Dynamics: Second Edition. Princeton University Press. https://doi.org/10.1515/9781400828722

Bosma, A. (1981). 21-cm line studies of spiral galaxies. II. The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types. The Astronomical Journal, 86, 1825. https://doi.org/10.1086/113063

Burkert, A. (1995). The Structure of Dark Matter Halos in Dwarf Galaxies. The Astrophysical Journal Letters, 447(1), L25–L28.

De Blok, W. J. G. (2010). The Core‐Cusp Problem. Advances in Astronomy, 2010(1), 789293. https://doi.org/10.1155/2010/789293

De Blok, W. J. G., & McGaugh, S. S. (1997). The dark and visible matter content of low surface brightness disc galaxies. Monthly Notices of the Royal Astronomical Society, 290(3), 533–552. https://doi.org/10.1093/mnras/290.3.533

Di Paolo, C., Salucci, P., & Erkurt, A. (2018). The universal rotation curve of low surface brightness galaxies IV: The interrelation between dark and luminous matter. https://doi.org/10.48550/ARXIV.1805.07165

Famaey, B., & McGaugh, S. S. (2012). Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions. Living Reviews in Relativity, 15(1), 10. https://doi.org/10.12942/lrr-2012-10

Freeman, K. C. (1970). On the Disks of Spiral and so Galaxies. The Astrophysical Journal, 160, 811. https://doi.org/10.1086/150474

Jiao, Y., Hammer, F., Wang, H., Wang, J., Amram, P., Chemin, L., & Yang, Y. (2023). Detection of the Keplerian decline in the Milky Way rotation curve. Astronomy & Astrophysics, 678, A208. https://doi.org/10.1051/0004-6361/202347513

Katz, H., Lelli, F., McGaugh, S. S., Di Cintio, A., Brook, C. B., & Schombert, J. M. (2016). Testing Feedback-Modified Dark Matter Haloes with Galaxy Rotation Curves:

Estimation of Halo Parameters and Consistency with $Λ$CDM. https://doi.org/10.48550/ARXIV.1605.05971

Karukes, E. V., Salucci, P., & Gentile, G. (2015). The dark matter distribution in the spiral NGC 3198 out to 0.22 Rvir. Astronomy & Astrophysics, 578, A13. https://doi.org/10.1051/0004-6361/201425339

Lelli, F., McGaugh, S. S., & Schombert, J. M. (2016a). SPARC: MASS MODELS FOR 175 DISK GALAXIES WITH SPITZER PHOTOMETRY AND ACCURATE ROTATION CURVES. The Astronomical Journal, 152(6), 157. https://doi.org/10.3847/0004-6256/152/6/157

Lelli, F., McGaugh, S. S., & Schombert, J. M. (2016b). THE SMALL SCATTER OF THE BARYONIC TULLY–FISHER RELATION. The Astrophysical Journal Letters, 816(1), L14. https://doi.org/10.3847/2041-8205/816/1/L14

Martinsson, T. P. K., Verheijen, M. A. W., Westfall, K. B., Bershady, M. A., Andersen, D. R., & Swaters, R. A. (2013). The DiskMass Survey: VII. The distribution of luminous and dark matter in spiral galaxies⋆. Astronomy & Astrophysics, 557, A131. https://doi.org/10.1051/0004-6361/201321390

McGaugh, S. S., Lelli, F., & Schombert, J. M. (2016). Radial Acceleration Relation in Rotationally Supported Galaxies. Physical Review Letters, 117(20), 201101. https://doi.org/10.1103/PhysRevLett.117.201101

Milgrom, M. (1983). A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. The Astrophysical Journal, 270, 365. https://doi.org/10.1086/161130

Moore, B., Quinn, T., Governato, F., Stadel, J., & Lake, G. (1999). Cold collapse and the core catastrophe. Monthly Notices of the Royal Astronomical Society, 310(4), 1147–1152. https://doi.org/10.1046/j.1365-8711.1999.03039.x

Navarro, J. F., Frenk, C. S., & White, S. D. M. (1996). The Structure of Cold Dark Matter Halos. The Astrophysical Journal, 462, 563. https://doi.org/10.1086/177173

Navarro, J. F., Frenk, C. S., & White, S. D. M. (1997). A Universal Density Profile from Hierarchical Clustering. The Astrophysical Journal, 490(2), 493–508. https://doi.org/10.1086/304888

Oh, S.-H., De Blok, W. J. G., Brinks, E., Walter, F., & Kennicutt, R. C. (2011). DARK AND LUMINOUS MATTER IN THINGS DWARF GALAXIES. The Astronomical Journal, 141(6), 193. https://doi.org/10.1088/0004-6256/141/6/193

Persic, M., Salucci, P., & Stel, F. (1996). The universal rotation curve of spiral galaxies—I. The dark matter connection. Monthly Notices of the Royal Astronomical Society, 281(1), 27–47. https://doi.org/10.1093/mnras/278.1.27

Primack, J. R. (2024). Galaxy Formation in ΛCDM Cosmology. Annual Review of Nuclear and Particle Science, 74(1), 173–206. https://doi.org/10.1146/annurev-nucl-102622-023052

Reid, M. J., Menten, K. M., Brunthaler, A., Zheng, X. W., Dame, T. M., Xu, Y., Wu, Y., Zhang, B., Sanna, A., Sato, M., Hachisuka, K., Choi, Y. K., Immer, K., Moscadelli, L., Rygl, K.

L. J., & Bartkiewicz, A. (2014). TRIGONOMETRIC PARALLAXES OF HIGH MASS STAR FORMING REGIONS: THE STRUCTURE AND KINEMATICS OF THE MILKY WAY. The Astrophysical Journal, 783(2), 130. https://doi.org/10.1088/0004-637X/783/2/130

Roberts, M. S., & Whitehurst, R. N. (1975). The rotation curve and geometry of M 31 at large galactocentric distances. The Astrophysical Journal, 201, 327. https://doi.org/10.1086/153889

Rubin, V. C., & Ford, W. K., Jr. (1970). Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. The Astrophysical Journal, 159, 379. https://doi.org/10.1086/150317

Rubin, V. C., Thonnard, N., & Ford, W. K., Jr. (1980). Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/. The Astrophysical Journal, 238, 471. https://doi.org/10.1086/158003

Salucci, P., & Burkert, A. (2000). Dark Matter Scaling Relations. The Astrophysical Journal, 537(1), L9–L12. https://doi.org/10.1086/312747

Sanders, R. H., & McGaugh, S. S. (2002). Modified Newtonian Dynamics as an Alternative to Dark Matter. Annual Review of Astronomy and Astrophysics, 40(1), 263–317. https://doi.org/10.1146/annurev.astro.40.060401.093923

Sofue, Y. (2012). Grand Rotation Curve and Dark-Matter Halo in the Milky Way Galaxy. Publications of the Astronomical Society of Japan, 64(4), 75. https://doi.org/10.1093/pasj/64.4.75

Sofue, Y. (2013). Mass Distribution and Rotation Curve in the Galaxy. In T. D. Oswalt & G. Gilmore (Eds.), Planets, Stars and Stellar Systems (pp. 985–1037). Springer Netherlands. https://doi.org/10.1007/978-94-007-5612-0_19

Sofue, Y., & Rubin, V. (2001). Rotation Curves of Spiral Galaxies. Annual Review of Astronomy and Astrophysics, 39(1), 137–174. https://doi.org/10.1146/annurev.astro.39.1.137

Van Albada, T. S., Bahcall, J. N., Begeman, K., & Sancisi, R. (1985). Distribution of dark matter in the spiral galaxy NGC 3198. The Astrophysical Journal, 295, 305. https://doi.org/10.1086/163375

Walter, F., Brinks, E., De Blok, W. J. G., Bigiel, F., Kennicutt, R. C., Thornley, M. D., & Leroy, A. (2008). THINGS: THE H I NEARBY GALAXY SURVEY. The Astronomical Journal, 136(6), 2563–2647. https://doi.org/10.1088/0004-6256/136/6/2563

Young, B.-L. (2017). A survey of dark matter and related topics in cosmology. Frontiers of Physics, 12(2), 121201. https://doi.org/10.1007/s11467-016-0583-4

Zhang, T. (2024). Dynamic Spacetime: Key to the Mysteries of Dark Matter and Dark Energy. Journal of Modern Physics, 15(04), 416–434. https://doi.org/10.4236/jmp.2024.154018

Zwicky, B., Meyling, J., & Appell, D. (2021). Fritz Zwicky and the earliest prediction of dark matter. Physics World, 34(5), 24–25i. https://doi.org/10.1088/2058-7058/34/05/28

Steel Decarbonisation and Cement Sustainability: A Literature Review of Slag-Based Industrial Symbiosis By Leon Langford

Abstract

Steel production is essential to modern infrastructure but remains a major source of greenhouse gas emissions, making its decarbonisation a significant industrial challenge. This literature review examines the main steel decarbonisation routes, the constraints affecting their deployment, and the changes in slag and co-product streams associated with different production pathways. Attention is given to the relationship between steel and cement, where ground granulated blast-furnace slag (GGBFS) can reduce clinker demand and associated cement emissions. A shift away from blast-furnace steelmaking may therefore reduce GGBFS availability, creating a trade-off between lower-carbon steel production and established cement decarbonisation strategies. The review also considers whether BOF, EAF and DRI-EAF slags could provide alternative cementitious or construction uses, while recognising differences in their chemistry, treatment requirements and performance. Overall, the literature indicates that no single steelmaking route is universally preferable, since outcomes depend on factors such as electricity supply, hydrogen availability, scrap quality, infrastructure and CCUS performance. Steel and cement should therefore be understood as linked material systems, with decarbonisation strategies assessed across both industries rather than in isolation.

1. Introduction

Steel is one of the main materials used by modern industry. It is required for buildings, railways, machinery and low-carbon infrastructure, including wind turbines and electricity transmission systems [1]. This creates a difficult sustainability problem: steel is necessary for decarbonisation, but steelmaking is itself a major source of greenhouse gas emissions. CO_2 refers specifically to carbon dioxide, whereas CO_2e expresses the warming effect of different greenhouse gases as an equivalent quantity of CO_2 [2], [3]. In 2024, global crude steel production was approximately 1.886 billion tonnes, with an average emissions intensity of around 2.18 tonnes of CO_2e per tonne of steel when Scope 1, 2 and 3 emissions are included [2]. This is equivalent to roughly 4.1 billion tonnes of CO_2e, or about 7-8% of global anthropogenic greenhouse gas emissions [2]. Blast-furnace-basic-oxygen-furnace production, or BF-BOF, still accounts for about 70% of global steel output, while announced near-zero-emission capacity remains below the level required by current transition pathways [4].

Existing studies generally compare steelmaking routes through emissions intensity, energy demand, technology maturity, infrastructure requirements and cost [1], [4], [5], [6], [7], [8], [9]. Work on cement and emissions definitions separately considers clinker reduction, supplementary cementitious materials and the treatment of co-products [1], [3]. Fewer studies bring these areas together in one transparent comparison. This literature review addresses that narrower gap by comparing the main steelmaking decarbonisation routes and examining how differences in their slag and co-product streams may affect cement decarbonisation. The

distinction between the routes is important because the main sources of emissions change depending on whether production relies mainly on iron ore, scrap, carbon, hydrogen or electricity.

The main ore-based route is BF-BOF steelmaking. Iron ore is reduced in a blast-furnace using coke, coal and carbon monoxide to produce hot metal, which is then refined in a basic oxygen furnace [1], [7]. BF-BOF with carbon capture, utilisation and storage, or CCUS, retains the same basic production route but captures carbon dioxide from selected process-gas streams [1], [7]. Direct reduced iron-electric arc furnace steelmaking, or DRI-EAF, reduces iron ore without first melting it and then melts the solid direct reduced iron in an electric arc furnace. Natural-gas DRI-EAF uses reducing gases produced from natural gas, whereas hydrogen DRI-EAF uses hydrogen as the main reducing agent [1], [6], [7]. Scrap-EAF mainly emulsifies recycled steel rather than producing new iron from ore. An electric arc furnace is therefore a furnace type rather than a single feedstock route, since it may be charged with scrap, direct reduced iron or a mixture of both [2], [1].

BF-BOF is difficult to decarbonise because carbon is used both as a source of heat and as a chemical reducing agent. Carbon monoxide removes oxygen from iron oxides to form metallic iron. A simplified equation is:

$$Fe_2O_3 + 3CO \rightarrow 2Fe + 3CO_2$$

This equation does not describe every reaction or process interaction inside a blast-furnace, but it shows why carbon-based reduction produces carbon dioxide as part of the chemistry [1], [7]. Emissions also arise from coke production, ore preparation, hot stoves, lime use, electricity, refining and casting [1], [7]. Efficiency improvements can lower energy demand, but they cannot remove the central process emissions while reduction remains carbon-based [1], [7]. Deep decarbonisation may therefore require changes to the reducing agent, the production route, supporting infrastructure and the treatment of residual emissions.

The comparison in this paper treats each route as a process pathway with different emission sources, energy requirements, feedstock constraints and co-product outputs. BF-BOF, BF-BOF with CCUS, natural-gas DRI-EAF, hydrogen DRI-EAF and Scrap-EAF will be assessed using published evidence on emissions, energy use, technology development and feedstock limitations [1], [4], [5], [6], [7], [8], [9]. Where emissions intensities are reported, values expressed per tonne of crude steel are treated as the most directly comparable because worldsteel data, IEA definitions and the ResponsibleSteel standard commonly use crude steel as the basis for emissions comparison [2], [3], [10]. This is more consistent than using finished steel, for which rolling, coating, product mix and manufacturing yield can change the reported result. Some process studies report results per tonne of liquid steel before casting, including Vogl et al. [8]. These values are therefore identified as liquid-steel results rather than being treated as directly equivalent to crude-steel values. This review also distinguishes whether each source

reports direct carbon dioxide emissions, scope 1 and 2 emissions, or broader carbon dioxide-equivalent emissions [2], [3], [10].

The practical constraints differ between routes. The hydrogen DRI-EAF case considered here uses electrolytic hydrogen produced by splitting water in an electrolyser. Its emissions depend strongly on the carbon intensity of the electricity used for both electrolysis and electric-arc-furnace operation [1], [8], [9]. Vogl et al. modelled a hydrogen-based route requiring approximately 3.48 MWh of electricity per tonne of liquid steel, with most of the electricity used by the electrolyser [8]. This is a result for one process design and should not be treated as a universal value. Rosner et al. showed that hydrogen price, heating strategy, gas recycle design and process integration can materially affect both cost and plant performance [9]. Hydrogen DRI-EAF also requires suitable iron ore and large additions of low-carbon electricity, hydrogen production and transport infrastructure [1], [7]. Scrap-EAF depends on scrap availability, scrap quality and the carbon intensity of electricity [1], [11], [12]. CCUS routes depend on capture performance, additional energy demand and access to carbon dioxide transport and permanent storage [1], [7]. Existing plant lifetimes and continuing investment in coal-based capacity may also delay the replacement of high-emission assets [1], [5], [7]. These differences mean that no single route can be ranked as universally superior without considering the conditions under which it is integrated within.

Material efficiency and recycling are related but not identical. Material efficiency means reducing the quantity of steel required to provide the same service. This may be achieved through improved design, less over-specification, higher manufacturing yield and longer product life [1], [11], [12]. Reuse keeps a component in service without remelting it, while remanufacturing restores a product or component so that it can be used again [11], [12]. Recycling uses collected scrap as an input for new steel production, usually through an electric furnace [11], [12]. These approaches can reduce demand for primary iron, but recycling alone cannot decarbonise the sector because scrap availability is constrained by earlier production levels, long product lifetimes, collection rates and sorting systems [1], [6], [11], [12]. Scrap can also contain residual elements such as copper or tin, which may restrict its use in some steel grades [1]. Mission Possible Partnership therefore includes demand reduction, material efficiency, greater scrap use and continued primary production within the same transition framework [6].

The main cement-related co-product considered in this paper is blast-furnace slag. Granulated blast-furnace slag can replace part of the clinker in blended cement, which can lower the emissions associated with cement production [1], [3]. In this paper, the processed material used as a supplementary cementitious material will be referred to as ground granulated blast-furnace slag, or GGBFS. Other ferrous slag streams will be kept separate, and the analysis will not assume that slag produced by electric-arc-furnace routes is an equivalent cementitious substitute unless the selected evidence supports that comparison.

A comparison based only on direct steel-production emissions may therefore miss a wider system effect. If BF-BOF production falls, the availability of blast-furnace slag used as a

clinker substitute may also fall [1]. The interpretation of route-level emissions can also depend on the accounting rules applied to co-products [3]. For that reason, the review considers route-level decarbonisation performance and co-product implications as related but distinct aspects of the steel transition. The steelmaking routes are compared using published evidence on their emissions, energy requirements, feedstock needs, infrastructural constrains and their co-product outputs. The review then examines how changes in BF-BOF and EAF production may affect the availability and use of cementitious co-products.

The primary aim of this literature review is to assess steel decarbonisation pathways while considering both steel-production performance and the implications of route changes for slag and cement-sector co-products.. The first research question asks what the main steel decarbonisation routes are, and which technical, material, energy and infrastructure constraints can affect their deployment. The second asks how these routes differ in the slag and other co-product streams they produce. The third asks what consequences reduced GGBFS availability could have for cement decarbonisation. The fourth question asks to what extent BOF, EAF and DRI-EAF slags can provide alternative cementitious or construction uses. IEA and ResponsibleSteel terminology is used where appropriate [3], [10]. The intended contribution is an integrated literature-based assessment of steel decarbonisation and cement sustainability that treats steel and cement as linked material systems.

2. Methods

This study used a targeted literature review to compare the selected steel decarbonisation routes and examine their possible co-product implications for the cement sector. The main routes considered were blast-furnace-basic oxygen furnace (BF-BOF), BF-BOF with carbon capture, utilisation and storage (CCUS), natural-gas direct reduced iron-electric arc furnace (NG-DRI-EAF), hydrogen direct reduced iron-electric arc furnace (H_2-DRI-EAF) and scrap-electric arc furnace (scrap-EAF). These routes were selected to cover the main current and lower-carbon steelmaking pathways discussed in the paper [1], [7], [8], [9].

DRI refers to direct reduced iron and EAF to an electric arc furnace. NG-DRI-EAF uses reducing gases produced from natural gas, whereas H_2-DRI-EAF uses hydrogen as the main reducing agent. Emerging routes, including electrochemical ironmaking and hydrogen plasma reduction were discussed where relevant but were not treated as core pathways because the available evidence is less mature and less directly comparable [13]. The review therefore focused on synthesising and comparing published evidence rather than constructing a lifecycle, techno-economic or plant-level simulation.

2.1 Literature source selection

The literature review searched for combinations of the route names with terms linked to emissions intensity, energy demand, electricity carbon intensity, hydrogen production, scrap share, CCUS, slag generation, GGBFS availability and cement substitution. The review was targeted rather than systematic. Its purpose was to identify evidence relevant to the comparative

assessment of steel decarbonisation routes and their steel-to-cement co-product relationships. It was not intended to produce an exhaustive catalogue of all published literature on steel decarbonisation.

Peer-reviewed papers were prioritised for process-level and technology-specific data, while agency reports and standards were mainly used for wider industry context, emissions definitions and comparison methods. The IEA technology roadmap was used for the main steelmaking route categories and wider decarbonisation pathways [1]. The IEA definitions for near-zero and low-emissions steel were used when considering emissions boundaries and the crude-steel basis of comparison [3], while ResponsibleSteel was used to support comparisons involving different scrap shares [10]. Fan and Friedmann were used for the wider comparison of low-carbon steelmaking technologies and CCUS options [7]. Vogl et al. and Rosner et al. provided process and energy information for hydrogen DRI [8], [9].

For the co-product part of the study, the GGBFS availability report was used to distinguish between local clinker substitution and wider supply-constrained effects [14]. Peer-reviewed studies on carbonated steel slags, slag carbonation, steel slag aggregate and EAF slag valorisation were used to assess whether alternative steelmaking slags could have possible uses in cement and concrete [15], [16], [17], [18], [19], [20].

2.2 Evidence extraction and comparison

A structured evidence-extraction table was used to organise information from the selected sources. Each entry recorded information from an individual source relevant to the comparison of steelmaking routes. Reported values were retained on the basis used by the original source so that differences in units, system boundaries and assumptions remained visible during comparison. The principle fields recorded in the evidence-extraction table are summarised below:

Table 1: Evidence extraction

Field type	Variables recorded
Source information	Source ID, year, page or section, source URL
Route information	Route, region
Reported value	Metric, reported value, or range where provided, unit, reported basis
Boundary information	Emissions or system boundary, emissions accounting category, whether upstream emissions were included
Co-product information	Co-product or slag type
Data quality information	Data type, confidence
Comparison information	Comparability notes, including differences in the functional unit, emissions boundary

Reported values and ranges were recorded only where they were provided directly by a source. Where a source reported only one value, it was retained as reported without constructing an additional range. This prevented unsupported values from being introduced to the review.

Confidence was assigned qualitatively as high, medium or low. High confidence was assigned where a value came from a peer-reviewed study, recognised agency or technical standard and where the route, unit, reported basis and emissions boundary were clear with little interpretation required. Medium confidence was used where the source was suitable, but a conversion, minor inference or some interpretation of the system boundary was required. Low confidence was used where substantial inference was required or important information on the unit, basis or boundary was unclear. Low-confidence evidence could remain in the extraction table for reference, but values that could not be compared reliably were not used to support direct route-to-route numerical comparisons.

Where the same numerical estimate appeared in more than one review or secondary source, the original source was used where it could be identified. This reduced the risk of treating repeated reporting of the same estimate as separate independent data points. A separate comparison table summarised route-level evidence that could be compared on a clearly stated basis. It recorded the original reported value and unit together with the relevant system boundary and a note explaining the extent to which the evidence was directly comparable

Values reported per tonne of crude steel were treated as the most directly comparable basis where available. This basis is used in major steel emissions frameworks and therefore provides a useful reference when comparing evidence across production routes [3], [10]. Values reported using another product basis were not automatically treated as equivalent to crude steel. They were discussed on their original basis unless the source itself provided sufficient information for a clearly comparable conversion. For example, Vogl et al. report electricity demand per tonne of liquid steel, so this value was kept on its original liquid-steel basis unless a suitable conversion to crude steel could be made from the available information [8].

CO_2 and CO_2e values were also kept separate. Results using different emissions boundaries, including direct process emissions, electricity-related indirect emissions or wider upstream emissions, were not averaged together. This avoided creating apparently precise route comparisons from values that measured different parts of the production system [3].

2.3 A comparative synthesis framework

The selected routes were compared using published evidence on emissions intensity, energy requirements, feedstock constraints, infrastructure requirements, technology maturity and co-product generation. Numerical values were compared only when the production basis and emissions boundary were sufficiently similar. Where published studies used incompatible functional units or system boundaries, the values were retained on their original basis and interpreted qualitatively rather than recalculated.

2.4 Slag and GGBFS co-product review

A separate part of the review examined how a reduction in BF-BOF production could affect the supply of ground granulated blast-furnace slag (GGBFS) available to the cement

sector. This part of the analysis used two interpretations so that local clinker substitution was kept separate from the wider effect of moving a constrained co-product between markets.

In the local substitution case, GGBFS was treated as a clinker substitute where its use was technically suitable. Under this case, greater use of GGBFS in a particular cement market could reduce the amount of clinker required in that market.

The constrained co-product case treated GGBFS as a supply-limited material that may already be used elsewhere. Under this interpretation, allocating more GGBFS to one cement application does not necessarily create an equal additional global emissions saving because the same material may be displaced from an existing use. This distinction was included because GGBFS supply is constrained and a large proportion of available blast-furnace slag is already utilised [14].

Blast-furnace slag, basic oxygen furnace slag and EAF slag were recorded separately. BOF and EAF slags were not assumed to be direct substitutes for GGBFS because their mineralogy, composition, treatment requirements and cementitious behaviour differ. A supplementary cementitious material role was included only where literature supported the particular slag and treatment route being considered, including carbonation or other processing where relevant [15], [16], [17], [18], [20]. Steel slag use as aggregate was treated separately from clinker or cement substitution because it represents a different construction application [19].

2.5 Scope and limitations

This paper is a literature review and is not a full lifecycle assessment, detailed techno-economic assessment, plant design or forecast of the global steel industry. No laboratory experiments or industrial measurements were carried out.

3. Current steel manufacturing routes and emissions sources

This section focuses on how the main steelmaking routes operate, where their emissions originate and which co-products they generate. It is limited to process configuration, emissions sources and co-product formation. Section 4 considers the decarbonisation role, advantages and deployment constraints of each pathway. The routes differ mainly in how metallic iron is obtained and how energy is supplied for reduction and melting. BF-BOF production reduces iron ore using carbon-based reducing agents, DRI routes reduce iron ore in the solid state using a reducing gas before melting it in an electric arc furnace, and scrap-EAF mainly uses metallic iron that has already been produced during an earlier steelmaking cycle [1], [7].

3.1 BF-BOF

The blast-furnace-basic oxygen furnace (BF-BOF) route begins with iron ore, metallurgical coal and a flux, usually limestone. Depending on the feedstock and plant design, iron ore may first be prepared through sintering or pelletisation, while metallurgical coal is converted into coke. Coke provides carbon and energy within the blast-furnace and also helps maintain the structure required for gas flow through the furnace [1], [7].

Iron ore, coke and fluxes are charged into the blast-furnace while heated air is introduced from the bottom. Carbon monoxide formed within the furnace acts as a reducing agent and removes oxygen from the iron oxides. A simplified overall reaction is:

$$Fe_2O_3(s) + 3CO(g) \rightarrow 2Fe(s) + 3CO_2(g)$$

The actual reduction occurs through several intermediate iron oxides rather than through a single reaction. The reduced iron eventually melts to form carbon-rich hot metal, while fluxes combine with gangue and other mineral matter to form molten blast-furnace slag [1], [7]. The hot metal is then transferred to a basic oxygen furnace. Oxygen is blown into the vessel to oxidise excess carbon and other unwanted elements. Scrap steel is commonly added to the BOF charge, partly to absorb heat from the highly exothermic oxidation reactions and to contribute metallic iron before the liquid steel is cast [1], [7].

Emissions occur across the integrated BF-BOF plant. Raw-material preparation and the production of coke require energy, while sintering and pelletisation can require fossil fuels. The blast-furnace is a major source because carbon is used both as an energy source and as part of the reduction chemistry. Further emissions arise from fuel combustion, flux production, BOF operation, electricity use and casting [1], [7].

Process gases from coke ovens, the blast-furnace and the BOF can be recovered and used for heat or electricity generation elsewhere in the integrated plant. This can improve energy efficiency, although the gases still contain carbon originating from fossil feedstocks. Their reuse can reduce extra fuel demand but does not remove the fossil carbon already present in the system [1]. The reported emissions intensity of BF-BOF also depends on the system boundary. A value covering only the blast-furnace cannot be compared directly with one that also includes coke-making, ore preparation, electricity or upstream inputs [3].

BF-BOF produces two distinct major slag streams. Blast-furnace slag forms during ironmaking, while BOF slag forms during steel refining. Blast-furnace slag can be rapidly granulated and then ground to produce GGBFS. BOF slag has a different composition and is not treated as an equivalent cementitious material in this study [14], [17].

3.2 Scrap-EAF

Scrap-EAF production mainly uses steel that has already been reduced from iron ore during an earlier production cycle. Recovered scrap is collected, sorted and prepared before being charged into an electric arc furnace. The term scrap-EAF describes a mainly scrap-based route rather than a requirement for a completely pure scrap charge. Depending on the required steel grade and the available feedstock, the metallic charge may also contain DRI, hot-briquetted iron or pig iron [1].

Electrical energy supplied through graphite electrodes produces the high temperatures required to melt the metallic charge. Oxygen, natural gas and carbon can also be used during EAF operation for heating, refining and process control. Once the charge has melted, its composition is adjusted, unwanted elements are oxidised where possible and alloying elements are added before the steel is tapped and cast [1], [13].

Because primary ore reduction is largely outside the scrap-EAF route, its emissions arise from a different set of sources. Electricity generation can make a large contribution depending on the electricity supply. Natural gas, carbon additions, graphite electrodes, lime, scrap preparation and other upstream inputs may also contribute. Reported emissions therefore depend on whether electricity-related and wider upstream emissions are included within the system boundary [1], [3].

Slag forms during EAF refining as oxidised elements combine with added fluxes. Its composition varies with the metallic charge, scrap chemistry, flux additions, degree of oxidation and refining practice. Scrap-EAF therefore produces EAF slag rather than blast-furnace slag, and this slag should not be assumed to have the same composition or cementitious behaviour as GGBFS [20].

3.3 NG-DRI-EAF

Natural-gas direct reduced iron-electric arc furnace (NG-DRI-EAF) steelmaking is an ore-based route that avoids the conventional blast-furnace. Instead of melting the iron during reduction, iron oxide is reduced in the solid state, commonly in a shaft furnace. Suitable iron ore pellets or lump ore pass through the furnace while a hot reducing gas flows through the bed [1], [7], [9].

In modern reformer-based NG-DRI designs, natural gas is converted into a reducing gas containing mainly hydrogen and carbon monoxide. An external catalytic reformer can use natural gas together with recycled process gas from the shaft furnace to generate this mixture. The exact gas composition depends on the plant and reformer design [9]. Steam reforming can be represented in simplified form by:

$$CH_4(g) + H_2O(g) \rightleftharpoons CO(g) + 3H_2(g)$$

Both hydrogen and carbon monoxide can then remove oxygen from the iron oxides:

$$Fe_2O_3(s) + 3CO(g) \rightarrow 2Fe(s) + 3CO_2(g)$$

$$Fe_2O_3(s) + 3H_2(g) \rightarrow 2Fe(s) + 3H_2O(g)$$

In practice, reduction occurs through intermediate oxide stages and depends on temperature, gas composition and ore properties. The product is direct reduced iron, or DRI, which remains largely solid and porous. It is then transferred to an EAF and melted and refined into steel [7], [9].

The Rosner et al. NG-DRI reference design shows how these process stages can be integrated. Natural-gas-derived syngas is used as the shaft-furnace reductant, while natural gas is also used in the reformer and in heat applications elsewhere in the integrated mill [9]. This means that the reductant, heat supply and EAF operation contribute through separate parts of the process. NG-DRI-EAF removes the conventional coke ovens and blast-furnace, but carbon dioxide is still produced through natural-gas reforming, carbon monoxide reduction and fuel combustion. Electricity is also required for the EAF and other equipment. Depending on the boundary, pellet production and upstream natural-gas emissions may also be included [7], [9].

The downstream EAF produces EAF-type slag rather than blast-furnace slag. Gangue and remaining oxides entering with the DRI combine with lime, dolomitic additions and oxidation

products during EAF operation. Its composition therefore depends on the DRI feed, flux additions and EAF refining conditions [20]. A shift towards NG-DRI-EAF changes the co-product stream as well as the ironmaking process.

3.4 H_2-DRI-EAF

Hydrogen direct reduced iron-electric arc furnace (H_2-DRI-EAF) uses a similar basic process configuration to gas-based DRI, but hydrogen replaces most or all of the carbon-containing reducing gas in the shaft furnace. The main simplified reduction reaction becomes

$$Fe_2O_3(s) + 3H_2(g) \rightarrow 2Fe(s) + 3H_2O(g)$$

The oxygen removed from the iron oxide therefore leaves mainly as water rather than carbon dioxide. Other parts of the steelmaking system can still produce emissions depending on the plant design and the energy supply [8], [9]. In the H_2-DRI-EAF case considered in this study, hydrogen is produced through water electrolysis. It is supplied to the DRI shaft furnace and the resulting DRI is then melted and refined in an EAF [8].

Vogl et al. modelled an integrated hydrogen-based process requiring approximately 3.48 MWh of electricity per tonne of liquid steel, with most of this electricity used to produce hydrogen through electrolysis [8]. This value applies to the process design assessed by Vogl et al. and is not treated as a universal H_2-DRI-EAF value. It is also reported per tonne of liquid steel and is retained on that basis in this study.

Hydrogen reduction is more complicated than replacing one reducing gas with another. Ma et al. describe interactions between chemical phase transformations, gas transport and mechanical changes in the iron oxide during reduction. Cracks, pores and other microstructural changes affect the movement of hydrogen and water through the material and can influence reduction kinetics [21].

Rosner et al. distinguish between replacing natural-gas-derived syngas in the shaft furnace and replacing natural gas used for heating elsewhere in the integrated steel mill [9]. In their model, replacing the shaft-furnace syngas with renewable hydrogen reduced direct on-site CO_2 emissions by 76.3%. Replacing the remaining natural-gas heat applications increased the direct reduction to 84.9% [9]. These figures apply to the specific plant configurations modelled by Rosner et al. and refer to direct CO_2 rather than total life-cycle CO_2e.

As with NG-DRI-EAF, the final steelmaking stage produces EAF-type slag rather than blast-furnace slag. A transition towards H_2-DRI-EAF therefore changes the co-product stream rather than producing a lower-carbon version of the same BF slag. Future ore-based H_2-DRI-EAF slags require separate assessment because their composition and treatment differ from conventional GGBFS [20].

Table 2: Comparison of the main steelmaking routes considered in this study

Route	Iron input	Main energy input	Main reductant	Main emission sources	Main co-product or slag
BF-BOF	Iron ore converted to hot metal, with some scrap commonly added in the BOF	Metallurgical coal and coke, recovered process gases, electricity	Mainly CO generated from carbon and coke	Iron-oxide reduction, coke-making, sintering or pelletisation, fuel combustion, lime, BOF operation and electricity	Blast-furnace slag and BOF slag. Granulated blast-furnace slag is the feedstock for GGBFS [14]
Scrap-EAF	Mainly recovered steel scrap, with possible DRI, HBI or pig iron additions	Mainly electricity, with oxygen, natural gas and carbon additions	No primary ore reductant in the scrap-based route, although carbon may be used during EAF operation	Electricity generation, natural gas and carbon additions, electrodes, lime and upstream scrap processing	EAF slag and other furnace residues. Slag composition depends on charge chemistry, fluxes and refining practice [20]
NG-DRI-EAF	Iron ore pellets or suitable lump ore converted to DRI	Natural gas and electricity	Reformed gas containing H_2 and CO	Natural-gas reforming and combustion, CO reduction, electricity, pellet production and upstream natural gas were included	EAF-type slag. Composition is influenced by DRI gangue, flux additions and EAF operation [20]
H_2-DRI-EAF	Suitable iron ore converted to H_2-DRI	Electricity for hydrogen production and EAF operation	H_2	Electricity and hydrogen production depending on the boundary, pellet production, EAF inputs, lime and remaining carbon use	Ore-based EAF-type slag rather than BF slag. Its possible cementitious use depends on composition and treatment [20]

4. Decarbonisation pathways for steel

Section 3 described the physical production routes, their emissions sources and their coproducts. This section considers how the main pathways can reduce emissions, the advantages they offer and the constraints that affect deployment. A technically low-emission route is not equally suitable for every plant or region because energy supply, raw materials, infrastructure and existing assets differ [1], [7].

4.1 Energy and material efficiency

Energy efficiency reduces the energy required to produce a given quantity of steel without necessarily changing the main production route. Measures include improved heat recovery, process control, equipment efficiency, use of recovered process gases and better integration between plant operations [1]. These measures can reduce fuel and electricity demand without requiring complete route replacement.

The remaining potential is limited. Modern blast-furnaces already operate relatively close to practical minimum energy requirements, and improving energy efficiency does not remove the carbon required for conventional iron-ore reduction [1]. Energy efficiency therefore supports decarbonisation but cannot by itself remove the main process emissions of BF-BOF production.

Material efficiency works through a different mechanism by reducing the quantity of new steel required to provide the same service. Higher manufacturing yields, longer product lifetimes, component reuse and more material-efficient design can reduce steel demand [1], [11], [12]. Milford et al. found that energy and emissions efficiency alone were insufficient to meet the reduction examined in their model, while material-efficiency measures provided further abatement [22].

Energy efficiency and material efficiency should therefore be treated separately. The first reduces the energy required for production. The second reduces the amount of production required. Both reduce pressure on the steel system but neither removes the need to lower emissions from the steel that is still produced.

4.2 Scrap use and circularity

Greater scrap use can reduce demand for primary iron and is therefore an important decarbonisation pathway. Its emissions performance depends strongly on the electricity supplying EAF production. With low-carbon electricity, scrap-based production can have low operational emissions. A carbon-intensive electricity supply increases the indirect contribution [1], [3].

The main constraint is that scrap is finite. Steel can remain in buildings, vehicles, machinery and infrastructure for decades before becoming available for recycling. Present scrap availability therefore reflects historical steel production as well as current collection and recovery rates. Even with high recycling rates, available scrap is not sufficient to replace all primary steel production [1], [6], [10], [11], [12].

Scrap quality also matters. Mixed scrap streams can contain residual elements such as copper and tin that are difficult to remove during conventional steelmaking and can restrict the grades that can be produced. Better collection, sorting and separation can improve scrap quality, while cleaner iron units such as DRI can be used where dilution is required [13], [23].

Scrap-EAF therefore has a major role in reducing emissions from secondary steel production, but greater scrap use cannot remove the need for lower-emission primary ironmaking. ResponsibleSteel also treats scrap supply as constrained and uses a scrap-variable approach when comparing emissions intensity [10].

4.3 Hydrogen DRI-EAF

H_2-DRI-EAF addresses primary steel production by replacing carbon-based ore reduction with hydrogen. Its main decarbonisation role is therefore the removal of much of the fossil-carbon chemistry associated with reducing iron ore [7], [8], [9].

Deployment depends on the supply of low-carbon hydrogen and electricity. Electrolytic hydrogen does not automatically produce low-emission steel if the electricity supplying the electrolyser is carbon intensive. The EAF and other plant operations also require electricity, so the emissions intensity and scale of the electricity supply remain important [8], [13]. Suitable ore is another constraint. DRI shaft furnaces generally require iron ore with properties that meet direct-reduction requirements, commonly including high-grade pellets. Large-scale deployment therefore requires sufficient DR-grade ore and pellet capacity [1], [7], [9].

Hydrogen supply also requires electrolysers, electricity generation, grid or dedicated power connections, hydrogen transport or on-site production, storage where required and new plant investment. These requirements can differ substantially between regions [1], [7]. Rosner et al. estimated break-even hydrogen procurement prices of \$1.70/kg H_2 when hydrogen was used for shaft-furnace reduction and \$1.63/kg H_2 when the remaining natural-gas heat applications

were also replaced [9]. These are study-specific values based on the reference plant, energy prices and process assumptions used by Rosner et al. They are not universal hydrogen-cost thresholds.

H_2-DRI-EAF can therefore provide deep reductions in emissions from new primary steel production, but its performance depends on low-carbon hydrogen, low-carbon electricity, suitable ore and the infrastructure required to supply them.

4.4 BF-BOF with CCUS

Carbon capture, utilisation and storage provides a different pathway because it can reduce emissions while retaining much of the existing BF-BOF production system. CO_2 is separated from selected process-gas streams, compressed and then transported for use or geological storage [1], [7]. This gives CCUS a possible role at existing integrated steelworks where complete route replacement would require major new investment. It can therefore reduce emissions from assets that may otherwise continue operating for a large part of their technical life [1], [7].

Capture rate and total route emissions reduction are not the same quantity. BF-BOF plants contain several emissions sources, so capturing a high proportion of CO_2 from one process stream does not mean that the same proportion of total plant emissions has been removed [7]. The coverage of capture across the plant therefore matters as much as the capture rate reported for an individual stream.

Utilisation and storage should also be separated. Geological storage is intended to retain captured CO_2 over the long term. Utilisation can place CO_2 into fuels, chemicals or other products, and some of these pathways can release the CO_2 later in the product chain. CO_2 utilisation should therefore not automatically be treated as permanent abatement [1].

CCUS requires infrastructure beyond the capture unit. Captured CO_2 must be compressed and transported, suitable storage sites must be available where permanent storage is intended, and stored CO_2 requires monitoring. Capture, compression and transport also require energy [1], [7]. BF-BOF with CCUS can reduce emissions from existing carbon-based assets, but its performance depends on capture coverage, residual emissions, the energy penalty and access to transport and storage infrastructure. It reduces emissions from the existing route rather than removing carbon from the underlying iron-reduction chemistry.

4.5 Biocarbon and methane pyrolysis

Biocarbon can replace part of the fossil carbon used in some iron and steelmaking applications and can support other decarbonisation pathways [7], [13]. Its use is limited by the physical and chemical functions required from carbon in steelmaking. Feedstock availability, land use, processing, transport and possible land-use change can also affect the upstream emissions associated with biomass [7], [13].

Methane pyrolysis provides a different supporting option by producing hydrogen from methane while separating carbon as a solid

$$CH_4(g) \rightarrow C(s) + 2H_2(g)$$

Bhaskar et al. assessed molten-metal methane pyrolysis as a hydrogen source for DRI [24]. The route still depends on natural gas, so upstream methane emissions remain relevant, while the solid carbon produced also requires a stable use or management route. Biocarbon and methane pyrolysis are therefore treated as complementary pathways rather than core pathways in the comparative review.

4.6 Electrochemical and plasma routes

Electrochemical ironmaking seeks to reduce iron compounds directly using electricity rather than first converting electricity into hydrogen. Proposed approaches include aqueous electrolysis and high-temperature molten oxide electrolysis [13]. These processes could provide alternative primary-iron routes where low-carbon electricity is available. Their present constraints include technology maturity, reactor scale-up, electrode and materials performance, reaction rates and cost [13]. They are therefore treated as possible future routes rather than as part of the main industrial route comparison.

Hydrogen plasma reduction uses reactive hydrogen species at high temperature and may combine reduction and melting more closely than conventional shaft-furnace DRI. Souza Filho et al. examined a hybrid process combining solid-state hydrogen reduction with hydrogen plasma reduction and melting [25]. The study demonstrates process potential but is based on experimental work rather than a complete commercial steelmaking route.

Electrochemical and plasma routes are therefore included qualitatively. Their future role depends on successful scale-up, industrial maturity and competitive operating costs [13], [25].

Table 3: Summary of steel decarbonisation pathways

Pathway	Main mechanism	Main strength	Main constraint	Role in this review
Scrap-EAF	Uses existing metallic iron in scrap and reduces the need for new ore reduction	Low energy demand relative to primary ore routes and low emissions where electricity is low carbon	Scrap quantity, scrap quality and electricity carbon intensity	Core pathway
H_2-DRI-EAF	Replaces carbon-based ore reduction with hydrogen reduction	Deep reduction of emissions from primary ironmaking where hydrogen and electricity are low carbon	Low-carbon electricity and hydrogen supply, DR-grade ore, infrastructure and cost	Core pathway
BF-BOF with CCUS	Captures CO_2 from selected emissions streams while retaining BFBOF production	Can reduce emissions from existing assets and may be retrofitted in some configurations	Capture coverage, residual emissions, energy penalty, CO_2 transport and permanent storage	Core pathway
Biocarbon and methane pyrolysis	Replaces part of fossil carbon or produces hydrogen while separating solid carbon	Can complement existing and emerging routes	Biomass supply and land use, upstream emissions, methane supply and solid carbon management	Core pathway
Electrochemical and plasma routes	Uses alternative iron-reduction processes based on direct electricity or reactive hydrogen plasma	Possible future route for low-emission primary ironmaking	Technology maturity, scale-up, materials performance and cost	Emerging pathway

Pathway	Main mechanism	Main strength	Main constraint	Role in this review
Energy efficiency	Reduces energy required by existing processes	Can reduce fuel and electricity demand without complete route replacement	Practical efficiency limits and continued process emissions	Supporting pathway
Material efficiency	Reduces the quantity of new steel required for the same service	Reduces the production required across the steel system	Depends on design, reuse, product lifetime and material-demand patterns	Supporting pathway

No single pathway is sufficient under every set of conditions. Scrap-EAF can reduce emissions where suitable scrap and low-carbon electricity are available, but scrap cannot supply all new steel demand. H_2-DRI-EAF can address primary steel production but requires low-carbon electricity, hydrogen infrastructure and suitable iron ore. BF-BOF with CCUS can reduce emissions from existing assets but retains carbon-based ironmaking and leaves residual emissions. Energy and material efficiency reduce the amount of energy and steel production required, while biocarbon, methane pyrolysis and electrochemical or plasma processes provide supporting or future options [1], [7], [13], [22], [24], [25].

The implication is that steel decarbonisation changes not only emissions intensity but also material flows. Moving away from BF-BOF reduces blast-furnace slag production and therefore the feedstock available for GGBFS. Greater EAF production creates alternative slag streams with different compositions and possible uses [14], [20]. Steel decarbonisation can therefore remove an established cementitious co-product at the same time as it creates new slag streams whose future use is less certain.

Section 5 examines whether these changes in steel production could alter the availability of cementitious co-products and whether alternative steelmaking slags can support, or constrain, wider cement-sector decarbonisation.

5. Cross industry implications of steel slag, GGBFS and cement sustainability

Steel decarbonisation affects more than the emissions produced during steelmaking. Changes in production routes also change the type and quantity of coproducts produced by the steel sector. This matters to cement because blast-furnace slag is already used to produce ground granulated blast-furnace slag (GGBFS), an established supplementary cementitious material. A shift away from BF-BOF production can reduce the supply of this material while increasing slag streams from EAF based steelmaking [14], [20], [26].

5.1 Clinker substitution and the difference between steel slags

Clinker is the main intermediate material used to produce conventional Portland cement. During clinker production, limestone and other raw materials are heated to high temperatures in a cement kiln. Carbon dioxide is released from fuel combustion and from the calcination of limestone. The main calcination reaction can be represented by

$CaCO_3(s) \rightarrow CaO(s) + CO_2(g)$

The decomposition of limestone accounts for a large part of the direct carbon dioxide emissions from clinker production, while fuel combustion produces most of the remaining direct emissions [27]. Reducing the amount of clinker required per tonne of cement can therefore

reduce emissions from both sources. The IEA identifies reducing the clinker to cement ratio as one of the main routes for lowering emissions from cement production, alongside energy efficiency, lower carbon fuels and carbon capture [27].

Supplementary cementitious materials, or SCMs, can replace part of the clinker in a cementitious system. GGBFS is an established SCM. It is produced from suitable blast-furnace slag that has been rapidly cooled to form a mainly glassy material and then ground. Its hydraulic properties allow it to replace part of the clinker used in cement and concrete [14], [26].

The term steel slag should not be used to imply that all slags have the same properties. Blast-furnace slag forms during ironmaking. BOF slag forms later when hot metal is refined into steel. EAF slag is produced during melting and refining in an electric arc furnace. These materials differ in chemical and mineralogical composition because they form from different feedstocks and under different process conditions [17], [18], [20], [26].

Table 4. Comparison of the main slag streams considered in this study

Material	Production route	Established cement use	Main advantages	Main concerns
Granulated blast-furnace slag, ground as GGBFS	BF ironmaking within BF-BOF	Yes. Established SCM after granulation and grinding [14], [26]	Reactive glassy structure, hydraulic properties and established clinker substitution	Global supply is constrained and falls as BF production declines. Grinding, processing and transport also carry environmental impacts
BOF slag	BOF steel refining	No established use equivalent to GGBFS. Treated and carbonated material is being studied as an SCM [15], [16], [17], [18], [26]	Calcium rich material with potential for carbonation, metal recovery and construction use after suitable treatment	Free CaO and MgO, iron content, volume stability, variable composition, leaching and additional treatment requirements
Scrap-EAF slag	Scrap based EAF melting and refining	Aggregate use is more established than SCM use [19], [20]	Existing construction uses and possible future valorisation after modification or treatment	Charge chemistry varies, while crystalline phases, iron oxides, volume stability and inconsistent reactivity can restrict cementitious use
Ore based DRI-EAF slag	EAF melting and refining of DRI or HBI	Emerging rather than established at industrial scale [20]	Can potentially be modified and granulated to produce a reactive material	Limited industrial data, different slag chemistry, processing requirements, quality control and scale up

This comparison shows why a change in steelmaking route cannot be represented by a single generic steel slag category. GGBFS has an established cement role, while BOF slag and EAF slags require separate assessment of chemistry, treatment and end use [14], [20], [26].

5.2 GGBFS availability and allocation

The benefit associated with GGBFS depends on the level at which its use is assessed. This creates two distinct perspectives. At a local product level, replacing clinker with GGBFS can lower the clinker content and embodied emissions associated with a particular cement or concrete product. The product can therefore show a lower reported emissions intensity because less new clinker is required [14], [27]. At a global constrained supply level, the same use may

not create an additional emissions reduction if the GGBFS would otherwise have displaced clinker in another market. Moving a limited quantity of GGBFS from one application to another changes its allocation but does not increase the total quantity available [14].

This distinction matters because GGBFS is already a constrained material. The review by the Institution of Structural Engineers and its partner organisations reports that around 90% of iron slag is already processed into GGBS each year and did not identify significant usable stockpiles that could provide a large additional supply [14]. The report therefore treats GGBS as a limited resource that is already close to full global utilisation [14]. This does not mean that existing GGBFS should no longer be used. Removing it from current applications could increase clinker demand elsewhere. The issue is whether increasing use in one project or region should automatically be counted as an additional global emissions saving when the same material already has another likely use.

The distinction becomes more important as steel production changes. GGBFS depends on the production of suitable blast-furnace slag. Lower BF-BOF production can therefore reduce the production base from which GGBFS is obtained. Krammer et al. identify declining blast-furnace slag availability as a consequence of the move towards greater EAF production, while Andersson et al. examine the different EAF slag streams expected as ore-based EAF steelmaking expands [20], [26].

Cement decarbonisation strategies that depend heavily on increasing GGBFS use may therefore become harder to sustain as BF production falls. Continuing carbon intensive blast-furnace ironmaking simply to maintain slag supply would not be a sensible response. The cement sector instead has to use available GGBFS carefully while developing other suitable SCMs and validating new slag streams created by lower emission steelmaking routes [14], [27], [20], [26].

5.3 Carbonated steel slag and future EAF slag

The expected reduction in blast-furnace slag supply increases interest in other steelmaking slags. BOF and EAF slags already have construction applications, particularly as aggregates, but their use as reactive SCMs is much less established than GGBFS [19], [20]. Wider cementitious use depends on controlling volume stability, reactivity, composition and environmental performance.

One limitation is the presence of free CaO and MgO in some steel slags. These phases can hydrate after the slag has been used and cause expansion. Wang et al. identify free CaO and MgO as major causes of soundness problems in untreated steel slag [17]. Untreated steel slag therefore cannot be assumed to behave in the same way as GGBFS.

Carbonation is one possible treatment. Carbon dioxide reacts with calcium and magnesium containing phases in the slag, converting some free CaO and MgO into more stable carbonate phases such as CaCO3 and MgCO3 [16], [17]. This can reduce expansion risk and store some carbon dioxide through mineralisation.

The amount of carbon dioxide taken up is not constant. Liu et al. report a wide range of uptake values across the studies in their review because performance depends on slag

composition, particle size, temperature, carbon dioxide concentration, reaction time and treatment method [16]. Mineral carbon uptake should therefore not be interpreted by itself as a net emissions saving. Grinding, transport, treatment energy, carbon dioxide supply and the operating conditions used for carbonation can all affect the overall environmental result. Biava et al. note that industrial application requires attention to process energy and wider life cycle impacts [18].

Carbonation can also affect environmental performance. Steel slags may contain elements whose leaching behaviour needs to be assessed before cement or concrete use. Carbonation can reduce the leaching of some metal ions by changing the phases in which they are contained, but the result depends on the original slag composition and treatment conditions [16], [18].

A further question is whether treated slag remains sufficiently reactive to act as an SCM. Srivastava et al. studied carbonated BOF slag and desulfurisation slag as partial cement replacements. In their experiments, 30 wt% of white Portland cement was replaced by carbonated slag. BOF slag carbonated at 60°C and 1 bar CO_2 produced compressive strength comparable with the reference cement under the conditions tested [15]. Greater degrees and rates of carbonation under other treatment conditions reduced later SCM reactivity [15]. Carbon uptake and cementitious performance therefore cannot be maximised independently. A treatment that increases carbonation does not necessarily produce the most reactive SCM. The final use depends on carbon uptake, mineral stability, hydraulic reactivity, durability and the required cement performance [15], [16], [17], [18].

EAF slag presents a related problem because its composition varies with the metallic feed, fluxes and operating conditions. Future ore-based EAF slags from DRI production may also differ from slags produced in conventional scrap-EAF plants [28]. Andersson et al. investigated an ore-based EAF slag from a furnace operating with hot briquetted iron. The slag was remelted, modified and water granulated before its reactivity was compared with commercial GGBFS [20].

Some modified slags achieved high amorphous contents and reactivities approaching those of GGBFS, while other compositions remained less reactive. Andersson et al. linked these differences to iron oxide content, slag basicity and the phases formed during cooling [20]. The study indicates that suitable EAF slag can potentially be developed into an SCM, but its chemistry and processing route have to be controlled.

Evidence for future H_2-DRI-EAF slag remains limited because the route does not yet have the industrial operating history of conventional BF-BOF production. Andersson et al. specifically notes the lack of industrial scale slag data for hydrogen based DRI-EAF production [20]. It would therefore be premature to assume that every tonne of blast-furnace slag lost through steel decarbonisation can be replaced by an equivalent tonne of EAF derived SCM.

Krammer et al. provide a more industry focused example of what would be required to move from laboratory potential towards a usable cement material. Their study used carbothermal reduction and rapid cooling to produce slag fractions from metallurgical residues, then assessed chemical composition, amorphous content, fineness, hydraulic activity and strength against

criteria drawn from EN 15167-1, Austrian standards and the literature [26]. A favourable laboratory strength result is therefore not sufficient on its own. A new slag derived SCM also requires repeatable production, controlled composition, suitable reactivity, environmental testing, quality control and a route to compliance with the standards required for commercial use [20], [26].

5.4 Industrial symbiosis between steel and cement

The use of steelmaking coproducts in cement is a form of industrial symbiosis. A material leaving one industrial process becomes an input to another. In the established BF-BOF system, blast-furnace slag produced during ironmaking can be granulated, ground and used as GGBFS to replace part of the clinker in cement. The steel sector gains an established outlet for a large coproduct stream, while the cement sector gains a material with useful hydraulic properties that can reduce clinker demand [14], [26].

This relationship means that the two sectors are partly linked through material supply. Steel producers do not operate blast-furnaces in order to make GGBFS. Slag availability follows the steelmaking route selected for steel production. When steelmaking moves away from BF-BOF towards scrap-EAF or DRI-EAF, the amount of conventional blast-furnace slag falls and different EAF slag streams increase [20], [26]. Steel decarbonisation therefore changes the inputs available to cement even when the decision is made entirely within the steel sector.

The loss of GGBFS does not mean that cement has no other decarbonisation options. Other suitable SCMs, including materials such as calcined clay, can reduce clinker demand, while treated BOF and EAF slags may provide further options where their performance is demonstrated [27], [26]. These materials are not automatic replacements for GGBFS. They differ in availability, chemistry, processing requirements, standards and cementitious behaviour.

Krammer et al. give an example of the processing required to turn this idea into an industrial material route. Their approach does not simply test whether untreated steel slag can be added to cement. It combines material selection, composition adjustment, carbothermal treatment, rapid cooling, metal recovery and testing against product requirements [26]. Both resulting streams have to be useful. The recovered metal needs a suitable route back into steel production, while the slag fraction needs the chemical and performance characteristics required for cement use.

For this relationship to operate at scale, a new slag derived SCM would need a sufficiently consistent supply, controlled processing, reliable quality, environmental compliance and a practical route between the steel plant and cement producer. Variability that can be tolerated in a disposal route may not be acceptable in a cement product. Positive laboratory strength results are therefore only an early stage in developing a reliable industrial material [20], [26].

The steel and cement sectors can respond to declining GGBFS through a combination of measures rather than a single replacement material. Existing GGBFS can continue to be used where it provides technical value. Cement producers can increase the use of other validated SCMs and reduce clinker demand through other established cement decarbonisation measures.

At the same time, BOF and EAF slags can be processed and assessed for cementitious use where their composition and treatment make this technically and environmentally suitable [14], [27], [20], [26].

Steel decarbonisation should therefore not be assessed only through emissions per tonne of steel. It also changes coproduct flows between industries. Moving away from BF-BOF can reduce access to an established cementitious material by lowering blast-furnace slag production. Growth in EAF based steelmaking can create new slag streams, but these require treatment, validation and quality control before they can support cement decarbonisation at scale [14], [20], [26]. The relationship between steel and cement is therefore a material flow problem as well as an emissions accounting problem. Decarbonising one sector changes the materials available to the other, so the wider effect depends on both the emissions reduction achieved in steel and what replaces the changing coproduct streams in cement.

6. Discussion

In the short term, increasing scrap use in electric furnaces is one of the more established options for reducing steelmaking emissions where sufficient suitable scrap and relatively low carbon electricity are available. Energy and material efficiency can also reduce emissions while lower-carbon production capacity is developed [1], [11], [12]. H_2-DRI-EAF is beginning to enter commercial deployment, although near-zero emissions capacity expected by 2030 remains limited. Some DRI projects may also begin operation using natural gas before later changing to hydrogen [4]. Over the medium term, H_2-DRI-EAF and CCUS could become more widely used as the required energy and infrastructure develops. In the longer term, the literature does not indicate that one route will replace all other forms of steelmaking. Scrap-EAF, hydrogen-based DRI and routes using CCUS are likely to have different roles depending on regional conditions [1], [7].

Regional conditions will strongly affect how this transition occurs. The IEA identifies energy prices, technology costs, raw material availability and regional policy as factors that affect the choice of steelmaking technology [1]. Electricity price and carbon intensity are especially important for EAF operation and electrolytic hydrogen production. Vogl et al. found that the production cost of H_2-DRI was strongly dependent on electricity price [8]. Scrap availability and quality also differ between regions because they depend on previous steel use, product lifetimes, collection and sorting systems [11], [12]. The availability of suitable iron ore can constrain DRI deployment, while hydrogen-based routes require large additions of electricity and hydrogen infrastructure. CCUS also depends on access to CO_2 transport and suitable permanent storage [1], [7]. A route that is practical in one region may therefore face much greater cost, feedstock or infrastructure constraints in another.

Common standards are needed to compare steel produced using different proportions of primary iron and scrap. The IEA near-zero emissions threshold changes with scrap share, ranging from 400 kg CO_2e per tonne of crude steel at the primary iron end of the scale to 50 kg CO_2e per tonne at the 100% scrap end [3]. ResponsibleSteel also uses a scrap-variable approach because

global scrap supply is limited and increasing scrap use at one site can reduce its availability elsewhere [10]. Procurement can support lower-emissions steel by creating demand and improving the case for investment in new production capacity. However procurement claims also require comparable emissions boundaries, reliable reporting and clear chain-of-custody rules. The IEA places physical emissions reductions at the centre of near-zero definitions, while clear accounting is needed to reduce the risk of the same emissions reduction being counted more than once [4], [3].

These pathways also involve several chemical engineering problems. Hydrogen reduction of iron oxides involves chemical reactions, transport processes and changes within the solid structure which can affect the overall reduction kinetics [21]. At plant level, H_2-DRI depends on heat supply and process integration. Rosner et al. showed that hydrogen pre-heating, top-gas recycle, recycle pressurisation and EAF off-gas use can affect plant energy demand and economics [9]. CCUS adds further process requirements because CO_2 must be captured from suitable process streams before transport and permanent storage [7]. Slag treatment presents a separate materials and process problem. Carbonation can reduce free CaO and MgO, improve volume stability and allow some treated steel slags to be considered for use as supplementary cementitious materials [15], [16], [17]. Steel slag used as aggregate requires suitable treatment and control of volume stability [19], while the potential SCM reactivity of EAF slag depends on its composition, phase formation and processing [20]. Material circularity therefore depends on how these co-products are treated and controlled rather than only on the amount of slag produced.

Steel and cement are also linked through these co-product flows. Lower BF-BOF production would reduce the blast-furnace slag available for conversion into GGBFS. Around 90% of iron slag is already processed into GGBS and the available material is considered a constrained resource that is almost fully utilised globally. Increasing its use in one market may therefore reduce its availability elsewhere rather than creating an equal additional global emissions reduction [14]. EAF slags cannot automatically replace GGBFS because their composition, phases and reactivity are different [20]. Carbonated steel slags may provide future opportunities as supplementary cementitious materials, while other steel slags can also be used in applications such as aggregate after suitable treatment [15], [16], [19]. Steel decarbonisation may therefore reduce the supply of an established cement-sector input while increasing other slag streams that require further treatment and validation. Development of EAF slag valorisation and other clinker-reduction materials could partly reduce the effect of declining GGBFS availability as steel production changes [14], [20].

6.1 Limitations

A limitation of the comparison is that the literature does not use one identical emissions boundary or product basis. Some sources report direct process CO_2, others report broader CO_2e values, and cradle-to-gate or other wider boundaries can include electricity-related and upstream emissions. These values are therefore not directly equivalent. The review kept CO_2 and CO_2e results separate, recorded the emissions boundary used by each source and only converted values

to the one-tonne-of-crude-steel basis where sufficient information was available for a clear conversion [3], [10]. Some process values could not be normalised without introducing additional assumptions. For example, Vogl et al. report electricity demand per tonne of liquid steel rather than per tonne of crude steel [8].

A further limitation concerns future slag routes. The evidence for future H_2-DRI-EAF slags and treated EAF-derived supplementary cementitious materials is much more limited than the evidence available for conventional GGBFS. Hydrogen-based DRI-EAF does not yet have the industrial operating history of BF-BOF, and Andersson et al. specifically identify the lack of industrial-scale slag data for hydrogen-based DRI-EAF production [20]. Existing studies therefore demonstrate technical potential rather than an established, quantified replacement pathway for GGBFS. Conclusions concerning future H_2-DRI-EAF slags and treated EAF-derived SCMs are consequently based on a qualitative review of the available evidence rather than a quantified forecast of future slag availability, quality or cement substitution.

7. Conclusions and recommendations

The key finding of this review is that steel decarbonisation cannot be evaluated only as a change in production emissions, it also changes co-product flows that affect connected industries such as cement. This study shows that there is no single steel decarbonisation pathway that is best under every set of conditions. Scrap-EAF can achieve large emissions reductions where sufficient high-quality scrap and low carbon electricity are available, but scrap supply and quality limit the amount of primary steel production that it can replace. H_2-DRI-EAF provides a route for reducing emissions from primary iron production, although its performance depends on low carbon electricity, low carbon hydrogen, suitable iron ore and supporting infrastructure. BF-BOF with CCUS may reduce emissions from existing integrated steelworks, but its effectiveness depends on capture coverage, residual emissions and access to suitable carbon transport and storage infrastructure [1], [7], [8], [9], [10].

Low carbon electricity and material circularity are therefore enabling conditions rather than optional additions to steel decarbonisation. Electrified steelmaking and electrolytic hydrogen depend directly on the emissions intensity of the electricity supply, while greater recovery, sorting, reuse and recycling of steel can reduce the amount of new primary iron that must be produced [1], [11], [22]. These measures need to develop alongside changes in the production route.

The comparison also shows the importance of using consistent emissions boundaries. Direct process CO_2, electricity related emissions, upstream emissions and wider CO_2e results should not be treated as equivalent unless the same parts of the production system are included. Scrap share and product basis should also be stated clearly when route emissions are compared [3], [10]. Co-products require similar care. Moving away from BF-BOF can reduce the supply of blast-furnace slag available for GGBFS while increasing EAF slag streams with different compositions and treatment requirements [14], [20], [26]. Steel decarbonisation should therefore be assessed through both emissions intensity and the material flows created by changes in production.

A full life cycle assessment could examine the wider consequences of route changes in greater detail. More industrial data are also needed for ore based DRI-EAF slags to determine whether these materials can provide reliable cementitious uses at scale [20]. Further experimental work could examine how slag composition, grinding and carbonation conditions can be optimised while maintaining reactivity, environmental performance and consistent product quality [15], [16], [17], [18], [20], [26]. These areas would help determine how steel decarbonisation can reduce emissions without overlooking its wider effects on connected material systems.

Works Cited

[1]“Iron and Steel Technology Roadmap – Analysis,” IEA. [Online]. Available: https://www.iea.org/reports/iron-and-steel-technology-roadmap

[2]“Climate change and the production of iron and steel - 2025,” worldsteel.org. [Online]. Available: https://worldsteel.org/climate-action/climate-change-and-the-production-of-iron-and-steel

[3]“Executive Summary – Definitions for Near-Zero and Low-Emissions Steel and Cement, and Underlying Emissions Measurement Methodologies – Analysis,” IEA. [Online]. Available: https://www.iea.org/reports/definitions-for-near-zero-and-low-emissions-steel-and-cement-and-underlying-emissions-measurement-methodologies/executive-summary

[4]“Steel – Breakthrough Agenda Report 2025 – Analysis,” IEA. [Online]. Available: https://www.iea.org/reports/breakthrough-agenda-report-2025/steel

[5]“Pedal to the Metal 2025 | Global Energy Monitor.” [Online]. Available: https://globalenergymonitor.org/research/pedal-metal-2025

[6]“Steel | materials - Mission Possible Partnership.” [Online]. Available: https://missionpossiblepartnership.org/steel-materials/

[7]“Low-carbon production of iron and steel: Technology options, economic assessment, and policy: Joule.” [Online]. Available: https://www.cell.com/joule/fulltext/S2542-4351(21)00095-7?_returnURL=https%3A%2F%2Flinkinghub.elsevier.com%2Fretrieve%2Fpii%2FS2542435121000957%3Fshowall%3Dtrue

[8] V. Vogl, M. Åhman, and L. J. Nilsson, “Assessment of hydrogen direct reduction for fossil-free steelmaking,” Journal of Cleaner Production, vol. 203, pp. 736–745, Dec. 2018, doi: 10.1016/j.jclepro.2018.08.279.

[9] “Green steel: design and cost analysis of hydrogen-based direct iron reduction† | Energy & Environmental Science | The Royal Society of Chemistry.” [Online]. Available: https://pubs.rsc.org/ee/article/16/10/4121/807435/Green-steel-design-and-cost-analysis-of-hydrogen

[10]“Steel Decarbonisation Scale | ResponsibleSteel.” [Online]. Available: https://www.responsiblesteel.org/steel-decarbonisation-scale

[11]“Towards a Circular Steel Industry.” [Online]. Available: https://www.irena.org/Publications/2023/Jul/Towards-a-circular-steel-industry

[12]“Circular economy policies for steel decarbonisation | OECD.” [Online]. Available: https://www.oecd.org/en/publications/circular-economy-policies-for-steel-decarbonisation_4cfb485d-en.html

[13]K. Rippy, R. T. Bell, and N. Leick, “Chemical and electrochemical pathways to low-carbon iron and steel,” npj Mater. Sustain., vol. 2, no. 1, p. 33, Oct. 2024, doi: 10.1038/s44296-024-00036-6.

[14]“The efficient use of GGBS in reducing global emissions - The Institution of Structural Engineers.” [Online]. Available: https://www.istructe.org/resources/guidance/efficient-use-of-ggbs-in-reducing-global-emissions/

[15]S. Srivastava, M. Cerutti, H. Nguyen, V. Carvelli, P. Kinnunen, and M. Illikainen, “Carbonated steel slags as supplementary cementitious materials: Reaction kinetics and

phase evolution," Cement and Concrete Composites, vol. 142, p. 105213, Sep. 2023, doi: 10.1016/j.cemconcomp.2023.105213.
[16]X. Liu, P. Wu, X. Liu, Z. Zhang, and X. Ai, "The Utilization of Carbonated Steel Slag as a Supplementary Cementitious Material in Cement," Materials, vol. 17, no. 18, p. 4574, Sep. 2024, doi: 10.3390/ma17184574.
[17]S. Wang et al., "A Review on the Carbonation of Steel Slag: Properties, Mechanism, and Application," Materials (Basel), vol. 17, no. 9, p. 2066, Apr. 2024, doi: 10.3390/ma17092066.
[18]G. Biava, L. E. Depero, and E. Bontempi, "Accelerated Carbonation of Steel Slag and Their Valorisation in Cement Products: A Review," Span. J. Soil Sci., vol. 14, p. 12908, Apr. 2024, doi: 10.3389/sjss.2024.12908.
[19]Z. Ren and D. Li, "27," Materials, vol. 16, no. 17, p. 5841, Jan. 2023, doi: 10.3390/ma16175841.
[20]"Insights into the Valorization of Electric Arc Furnace Slags as Supplementary Cementitious Materials | Journal of Sustainable Metallurgy | Springer Nature Link." [Online]. Available: https://link.springer.com/article/10.1007/s40831-023-00778-y
[21]Y. Ma et al., "Hierarchical nature of hydrogen-based direct reduction of iron oxides," Scripta Materialia, vol. 213, p. 114571, May 2022, doi: 10.1016/j.scriptamat.2022.114571.
[22]R. L. Milford, S. Pauliuk, J. M. Allwood, and D. B. Müller, "The Roles of Energy and Material Efficiency in Meeting Steel Industry CO2 Targets," Environ. Sci. Technol., vol. 47, no. 7, pp. 3455–3462, Mar. 2013, doi: 10.1021/es3031424.
[23]D. Raabe, C. C. Tasan, and E. A. Olivetti, "Strategies for improving the sustainability of structural metals," Nature, vol. 575, no. 7781, pp. 64–74, Nov. 2019, doi: 10.1038/s41586-019-1702-5.
[24]A. Bhaskar, M. Assadi, and H. N. Somehsaraei, "Can methane pyrolysis based hydrogen production lead to the decarbonisation of iron and steel industry?," Energy Conversion and Management: X, vol. 10, p. 100079, Jun. 2021, doi: 10.1016/j.ecmx.2021.100079.
[25]I. R. Souza Filho et al., "Green steel at its crossroads: Hybrid hydrogen-based reduction of iron ores," Journal of Cleaner Production, vol. 340, p. 130805, Mar. 2022, doi: 10.1016/j.jclepro.2022.130805.
[26]"Valorisation of metallurgical residues via carbothermal reduction: A circular economy approach in the cement and iron and steel industry - Anna Christine Krammer, Klaus Doschek-Held, Florian Roman Steindl, Katharina Weisser, Christoph Gatschlhofer, Joachim Juhart, Dominik Wohlmuth, Christoph Sorger, 2024." [Online]. Available: https://journals.sagepub.com/doi/10.1177/0734242X241240040
[27]"Technology Roadmap - Low-Carbon Transition in the Cement Industry – Analysis," IEA. [Online]. Available: https://www.iea.org/reports/technology-roadmap-low-carbon-transition-in-the-cement-industry

Influential Factors of Effective Encapsulation Design in Fragrance Delivery By Liana Kim

ABSTRACT

Encapsulation methods of fragrance molecules and their effects along with different factors that influence fragrance delivery in encapsulation were investigated to further optimize the delivery of fragrances to satisfy the consumer desires of longer lasting scents/ top notes. Studies displayed that the control release of fragrance was affected by dilution, aging, concentration of perfume, humidity, release mechanisms, and the hedonic characteristics of fragrance molecules.

INTRODUCTION

Perfume which is a combination of three primary note levels– top, middle and base– are composed of semi volatile compounds. Various studies have been conducted on fragrance volatility and how to optimize the notes of fragrances to satisfy the consumers' desires. However, despite their demand for a longer lasting top note which is the lightest layer, it has been difficult to preserve such volatile compounds. Encapsulation, although it has a reputation of being costly, is an effective delivery method in controlling volatility of specific notes and enhancing release as it preserves volatile compounds in storage and delivery, protecting them from evaporation and oxidation (Quellet, 2001). Studies conducted on the structures of fragrance molecules have further expanded on the factors that influence volatility and fragrance delivery for a more cost-effective encapsulation design.

Methodology, Criteria of Systematic Review & Limitations

The collection of studies were performed so that it displays the factors and its effects that can be considered during designing the ideal encapsulation method. These factors and effects are linked to volatility as this can create a criteria of what components can be controlled for better encapsulation and allow controlled release. The updated PRISMA Flow Diagram was followed during the process of narrowing down reports as shown in Figure 1a). Excluded studies include the specific methods of encapsulation in the industry (3 reports), the physiological aspect and behaviors of fragrance molecules that shifts volatility and any factors in relation to the atmosphere such as the interaction between the encapsulated fragrance molecules and the surfactants in the air. A study focusing explicitly on flavor encapsulation and an irrelevant study was excluded during assessment of eligibility. Due to limitations in available resources and studies focusing merely on general encapsulation design in fragrance delivery, this review relies heavily on the findings from Mirzamani, 2024, supplemented by Quellet, 2001.

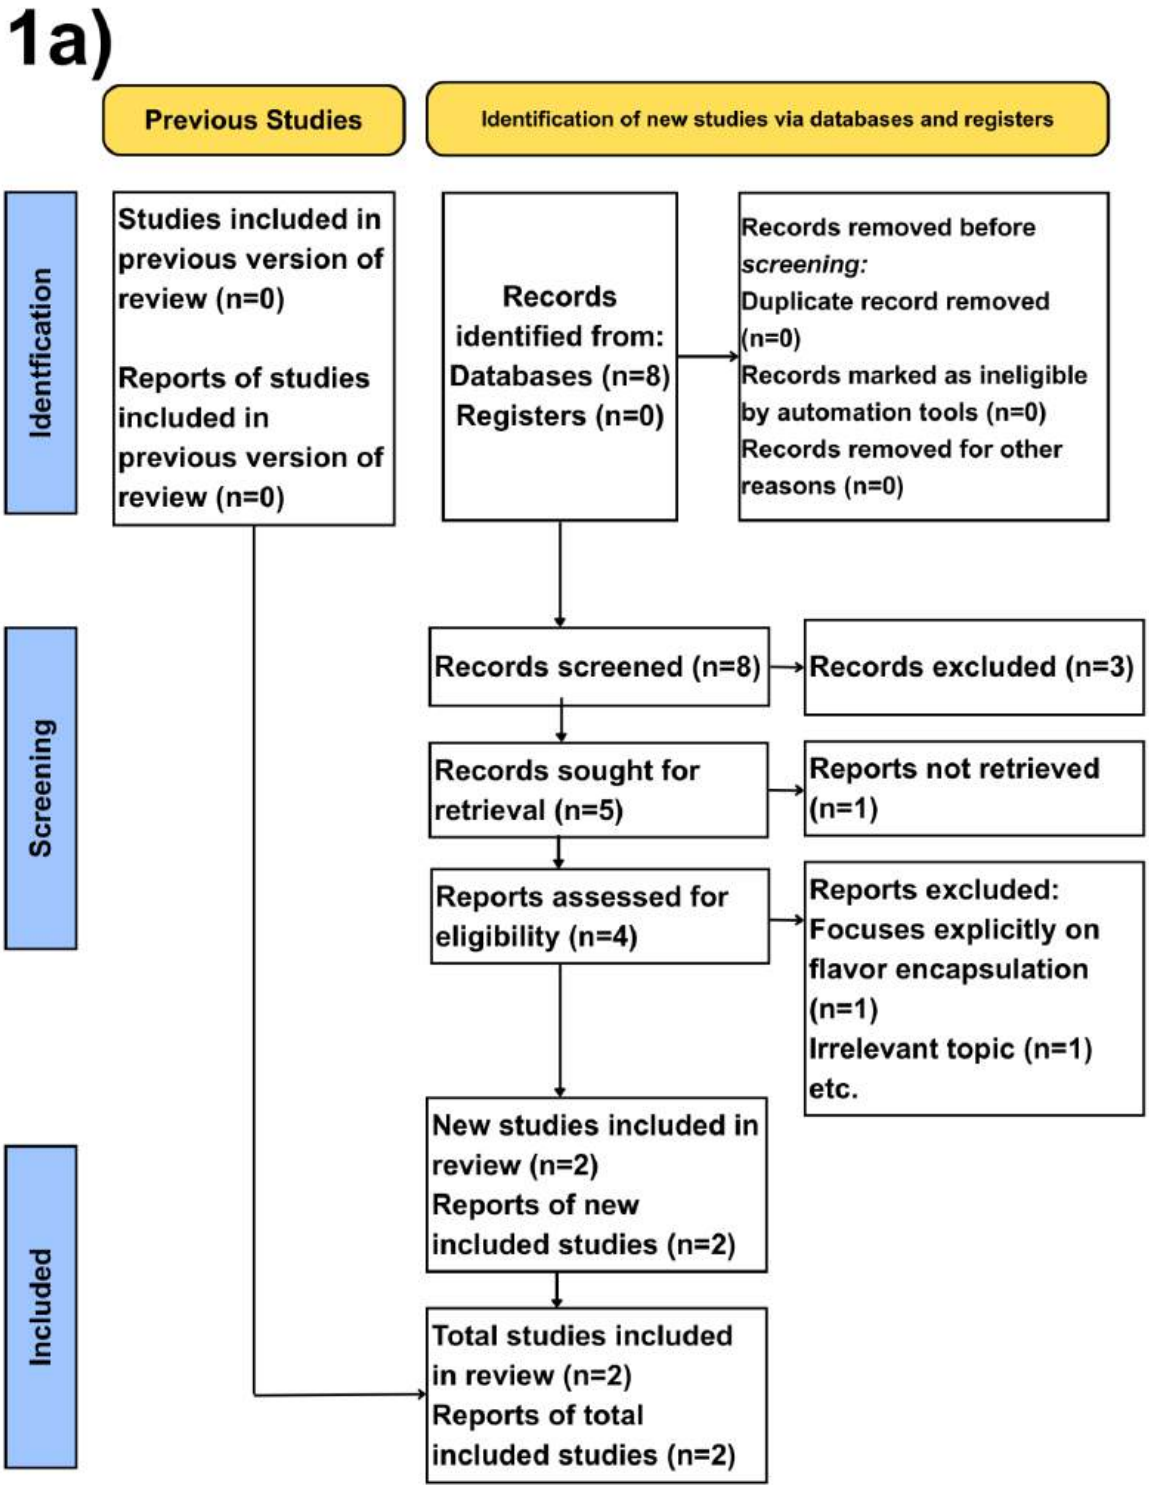


Figure 1a) Page MJ, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ 2021;372:n71. doi: 10.1136/bmj.n71

EXPERIMENTAL METHODOLOGIES

Dependent Variables Investigated in Dilution, Aging, and Concentration of Perfume

To analyze the experimental method of Mirzamani (2024), ST2S was used to represent branched-tail primary surfactant systems and SLE1S was used to represent the linear-tail primary surfactant system (both systems contained the cosurfactant CAPB) in the studies of dilution, aging, and concentration of perfume in relation to fragrance delivery. SANS data was then collected on NGB-SANS beamline and each data set was reduced in Igor Pro to avoid experimental errors (Mirzamani, 2024). The data were set to absolute scale and was circularly averaged to obtain 1D data and combined to form one data set. The data were then fitted with a smeared Uniform Ellipsoid form factor with Screened-Coulomb structure factor model using SasView version 5.0.3 (Mirzamani, 2024). The Uniform Ellipsoid provided information about the geometry of the micelles. $R_{a,}$ was the axis of rotation, R_b was the equatorial radius. Semi-major and semi-minor axes depended on which axis was greater. $R_a > R_b$ indicated prolate/ melon-shaped whereas $R_b > R_a$ indicated oblate/ flattened sphere (Mirzamani, 2024). The geometries and the semi-major and semi-minor axes were shifted upon the refinement of different variables.

Methodology for investigation of Aging and Surfactant Concentration

Samples of the surfactant system were prepared and equilibrated in glass vials at room temperature for 3 days prior to data collection for stability study in relation to aging (Mirzamani, 2024). GC-MS data were obtained using Agilent Model 6890 gas Chromatograph (GC). Clean 20mL headspace vials containing samples and magnetic stir bar were closed with PTFE septum caps and were set to equilibrate, then stirred and absorbed onto SPME fiber. Analysis to identify the PRMs were run on the SPME fiber after thermal desorption (Mirzamani, 2024). The changes in headspace of the surfactant systems were observed to determine the effects of various determinants in relation to fragrance delivery.

RESULTS & DISCUSSION

Aging

GC-MS studies exhibited the lower headspace concentration of older surfactant samples (Mirzamani, 2024), presenting the underlying mechanics behind evaporation of perfumes on the skin over time.

Dilution

The thermodynamic effect of dilution was studied to investigate its effect on the micelle shapes and behavior in relation to the surfactant systems. SANS data displayed that as dilution level went up, micelle volume decreased for both geometries of surfactant. When diluted to full strength, the length of semi-axes changed and made the ellipsoid less eccentric (Mirzamani,

2024). R_a decreased and R_b increased for the linear tailed system while the opposite was observed for the branched tailed system. While the linear-tailed system maintained its original oblate geometry, the branched-tailed system's geometry switched from prolate to oblate after dilution (Mirzamani, 2024). As shown in Figure 2a) unlike other factors which had consistent relationships with volatility, upon dilution, the headspace concentration dropped significantly after the first dilution and further decreased by smaller amounts after additional dilutions for both systems (Mirzamani, 2024). This result exhibits the perfume lost by dissolution in water which exhibited the reason behind the weakening of the strength of perfumes when diluted (Quellet, 2001). Thus, after rinsing, the volatility of fragrance molecules are affected in a way that some evaporate more/ less rapidly, depending heavily on their individual volatility and affinity with respect to the substrate or co-absorbed species (Quellet, 2001). Moreover, the difference in the loss of perfume concentrations is also attributed to the characteristic of the formulation of heavier fragrance compounds being the bottom notes; the heaviest which remains on the surface for a longer period of time (Quellet, 2001).

Perfume & Surfactant Concentration

At lower surfactant concentrations, changes to branched-tail systems are observable whereas in higher surfactant concentrations, the samples become two-phase according to the SANS data. In extremely high surfactant concentration, a clear bottom layer and an opaque upper layer was observed, showing lamellar structures and charged ellipsoidal structures. As a result, both major axes in higher surfactant systems decreased due to high surfactant concentrations (Mirzamani, 2024). This was attributed to the destability caused by the formation of new lamellar structures from high surfactant concentrations which was the system's reaction with respect to the stress of the addition of perfume (Mirzamani, 2024). That is to say, lower surfactant systems were ideal as they efficiently solubilized the perfume and prevented phase separation, therefore decreasing volatility and preserving scent molecules. However, although the lower surfactant system is the most sufficient regarding prevention of phase separation, it isn't sufficient to create stronger perfume-surfactant interactions to decrease the perfume's activity, which is the limitation of the ideal control release system. In relation to surfactant geometries, the linear- tail mixed-surfactant micelles were better at accommodating additional perfume as they didn't undergo changes in micellar geometry and phase-separation. In addition, GC-MS results also demonstrated that although headspace concentrations increased for both, the headspace concentration didn't rise proportionally to the overall perfume concentration which indicates that more perfume accord or specific PRM is necessary to achieve greater fragrance strength in delivery (Mirzamani, 2024).

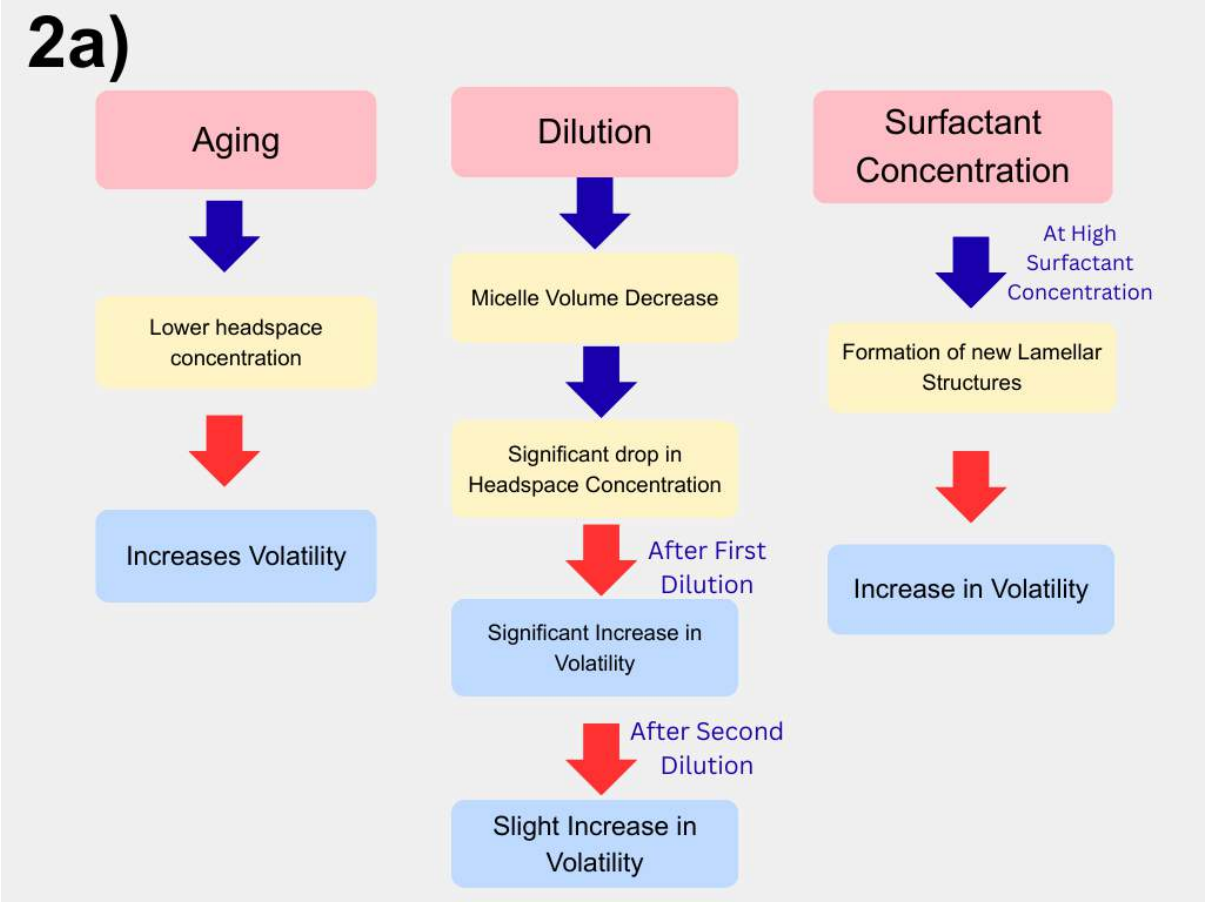


Figure 2a) Flow chart of Factors Studied in relation to Volatility

Factors studied by Mirzamani(2024) all had a direct relationship with volatility as increasing such factors either led to lower headspace concentration or led to formation of new structures, eventually leading to increase in volatility. However, dilution was the only factor which exhibited different results upon each trial; micelle volume decreased for all trials but the amount that volatility increased by differed as increased dilution after the first trial only caused a slight increase in volatility.

Absence vs Presence of Water and Humidity

Water has the effect of plasticizing, swelling and dissolving matrices which results in poor retention and decreases the ability of control in release mechanisms (Quellet, 2001). According to Quellet (2001), for compounds with less than 80% humidity and low water activity, hydrophilic matrices are effective in storage as it has low permeability to oxygen whereas for aqueous compounds with high relative humidity, hydrophobic carriers are more effective although they are permeable to both hydrophobic substances and oxygen and has limited capability in storage.

Release Mechanisms

The type of release mechanism is also a prime determinant when it comes to encapsulation design. For instance, burst-like release mechanisms are triggered by mechanical stresses and application of water (Quellet, 2001). Controlling the transport kinetics for ideal release of fragrance is essential for the production of sustained release (Quellet, 2001).

Hedonic Character

The hedonic characteristic of fragrance molecules– the subjective perception in scents– adds another challenge in fragrance delivery design. This is because the delivery system must be olfactory neutral and selectivity of both retention and permeability must be as low as possible to prevent composition distortion during use and preserve the original scent for a prolonged period of time (Quellet, 2001).

CONCLUSION

For a cost-effective design of encapsulation, it is important to find the appropriate middle point of the quality and cost of the design. As shown in Figure 2b) when forming an optimal encapsulation for fragrance molecules, factors that should be considered include: aging, dilution, perfume concentration, humidity, the appropriate release mechanisms and hedonic character. Although it is nearly impossible to control aging to increase the preservation rate of notes of perfumery, it is crucial to envision other determinants. Dilution decreases the overall strength of the fragrance, making dilute encapsulations perfect for rinse-off products (Mirzamani, 2024; Quellet, 2001). In future research, hydrophobic encapsulation materials are to be researched for better diversity in options. As the findings in relation to surfactant concentration to the stress of additional perfume concentration, lower surfactant concentration systems were in fact better capable in preventing phase separation despite causing weaker strength of fragrance. For future encapsulation design, the maximum point where phase separation can be prevented but emit greater strength of fragrance should be further investigated for consumer products that are meant to emit stronger odors. Furthermore, release mechanisms should be chosen according to the product prototype as undesired emissions can be prevented during storage depending on the purpose of the product. For instance, laundry detergents which utilizes heat as its trigger mechanism preserves its aroma during room temperature storage as the encapsulation is designed to be unaffected by standard temperatures (Quellet, 2001).

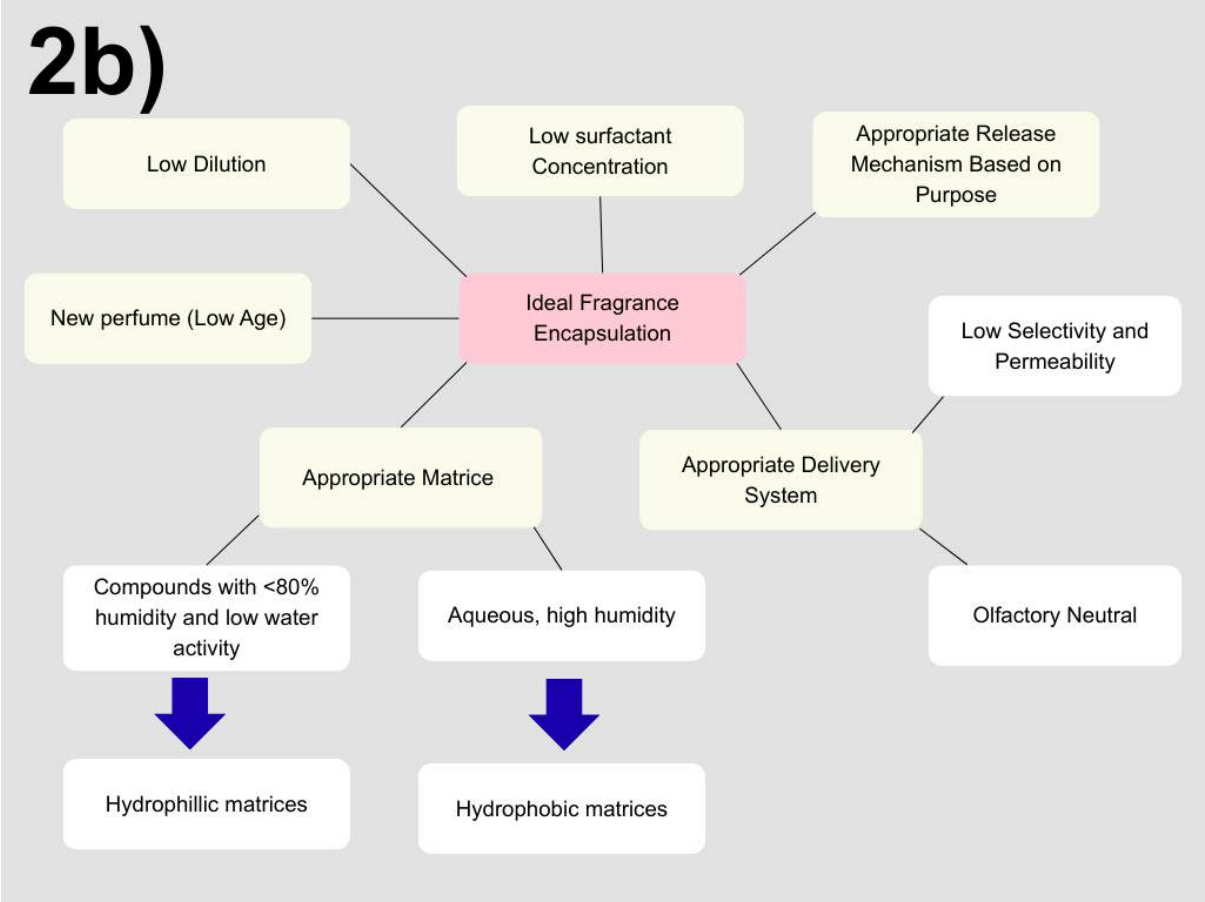


Figure 2b) Flow chart of Creating the Ideal Fragrance Encapsulation

Along with Figure 2a where the relationship between different factors and volatility was shown, this provides an overall flow chart of how much of which factor is ideal and what to consider in choosing the most appropriate fragrance encapsulation method.

Works Cited

Mirzamani M, Flickinger M, Jones RL, Ananthapadmanabhan K, Smith E, Kumari H. Effects of dilution and age on perfume release from two mixed surfactant systems containing hydrotropes. AIChE J. 2024; e18557. doi:10.1002/aic.1855

C. Quellet, M. Schudel, R. Ringgenberg, *Chimia* 2001, *55*, 421, DOI: 10.2533/chimia.2001.421.

REFERENCES FOR FIGURES

Figure 1a: Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ. 2021;372:n71. doi: 10.1136/bmj.n71.

Drag Force on a Wigley Hull: Equation vs. Simulation Across Velocities and Drafts

By Rithikk Vimal

KEYWORDS

Fluid dynamics, Drag force, Wigley hull, Resistance, Simulation

ABSTRACT

This study investigates the effect of varying hull velocities and draft lengths on the total resistance acting against a Wigley hull by comparing an analytical equation model with a digital simulation model. The results indicate that both velocity and draft length significantly increase total resistance, though differences in magnitude highlight the limitations of simplified analytical equations compared to comprehensive simulations.

SUMMARY

The study of fluid dynamics involves analyzing how physical factors such as velocity, gravitational strength, object length, and fluid viscosity influence the movement of objects through fluids. Accurate estimation of boat resistance is vital for marine engineering, yet theoretical equation models often struggle to balance simplicity and precision without complex integrals. The purpose of this research was to determine what is the effect of varying hull velocities (5, 10, 15, 25, 35, 45, 55, 65 m s-1) and Wigley hull draft lengths (0.020, 0.050, 0.150, 0.200, 0.210, 0.220, 0.240, 0.245 m) on the total resistance force acting against the boat when comparing the reduction of an equation and simulating as means of calculation. We hypothesized that as the velocity and draft increased, it would lead to the increase of total resistance nonlinearly. To test this, total resistance was evaluated across 64 unique draft-velocity pairings using an analytical equation model derived via the Buckingham Pi Theorem alongside ITTC frictional coefficients, and a digital simulation model via the **Prelimina Resistance Calculator**. The results demonstrated that total resistance increased nonlinearly with both velocity and draft length across both calculation methods, though the equation model yielded significantly lower absolute magnitudes due to the omission of complex wave-resistance components like Mitchell's integral. These findings conclude that while simplified equation models successfully capture qualitative nonlinear trends, digital simulations provide necessary precision for absolute resistance estimation.

INTRODUCTION

The study of fluids is considered fluid mechanics; more specifically, the movement of such objects in the fluids is the study of fluid dynamics. In fluid dynamics, the velocity, the pressure, the gravitational strength, the length of the object, and viscosity of water influence the object’s movement. For this research the deviations in values will be constrained to the velocity and length, to recognize differences in the Froude’s (Fr) and Reynold’s (Re) numbers (Nasa 2025). Using these constants, the research will attempt to appropriate resistance faced by a

specific model of hull, specifically the Wigley, through the usage of an equation model contrasting with simulation model. Although this research will not directly derive the perfect resistance through the equation model, it will attempt to estimate a similar, simpler equation in which the fluid constants: Reynolds and Froudes numbers are utilized.

The research will measure the effect of varying hull velocities (5, 10, 15, 25, 35, 45, 55, 65 m s-1) and Wigley hull lengths above water (0.020, 0.050, 0.150, 0.200, 0.210, 0.220, 0.240, 0.245 m) on the total resistance force acting against the boat using both an equation model and a simulation model to determine such value. The equation model will begin with the derivation of both the fluid constants through the Buckingham Pi Theorem. Following this, the ITTC coefficient of friction (for the viscous resistance) along with a similar Fr sinusoidal function (for the wave-loss resistance), which through nondimensionalization, were summed to create the equation model. Conversely, the simulation model involves the resistance calculator provided by Prelimina. These methods will vary through eight increments of hull draft lengths and the hull velocities, producing a total of 64 unique draft-velocity pairing for which the resistance will be measured. Following this process, a statistical analysis will be performed calculating the average total resistance (Avg R_T) for each pairing, along with standard deviation, and independently recognizing each independent variable's influence on the total resistance (R_T), producing models based on the data.

RESULTS

Resistance Calculations Across Draft-Velocity Pairings

The total resistance values obtained from the equation model and the simulation model across all 64 unique velocity-draft pairings revealed distinct numerical patterns. In the equation model data, values gathered at 5 m s-1 were observed to be higher than those gathered at 10 m s-1 and 15 m s-1 for most draft lengths except for the 0.020 m draft. Across the simulation model, total resistance scaled rapidly with velocity, reaching up to 10411 N at maximum velocity and draft parameters.

Statistical Determination and Regression Models

To analyze the aggregate impact of the parameters, average total resistance and standard deviations were calculated across varying draft lengths and hull velocities. For the equation model with varying draft lengths, the average total resistance progressed from 3.383 N at a 0.020 m draft to 286.877 N at a 0.045 m draft, accompanied by standard deviations reflecting distribution spread. When evaluating hull velocities independently in the equation model, averages ranged from 20.149 N at 10 m s-1 to 485.677 N at 65 m s-1.

Percentage error comparisons between the equation model and the simulation model indicated a magnitude variance ranging between 90% and 99% across respective pairings, demonstrating that the simplified equation model underestimated absolute resistance magnitudes compared to the simulation software. Regression analyses identified quadratic and exponential

models as the best fits for the data. For the equation model, draft length versus average total resistance fit a negative quadratic model yielding $y = -22.3 + 88.8x - 6.36x^2$ with a coefficient of determination of $R^2 = 0.957$. Hull velocity versus average total resistance in the equation model followed a positive quadratic model represented by $\$y = 54.6 - 31.2x + 13.5x^2$ with an $R^2 = 0.995$. Conversely, the simulation model data for velocity versus average total resistance fit an exponential curve of $y = 285e\{0.478x\}$ ($R^2 = 0.971$), while simulation draft length versus average total resistance followed a quadratic model of $y = 48.2 + 1250x - 106x^2$ ($R^2 = 0.947$).

DISCUSSION

The essay's aim was to see the effect of varying hull velocities (5, 10, 15, 25, 35, 45, 55, 65 m s-1) and Wigley hull drafts lengths (0.020, 0.050, 0.150, 0.200, 0.210, 0.220, 0.240, 0.245 m) on the total resistance acting against the boat when comparing the equation and simulation models, which yielded that as the draft lengths and hull velocities increased, it increased the total resistance experienced by the hull.

The simulation model's larger magnitude can be attributed to the fact that the equation model doesn't account for all aspects of the resistance acting upon the hull. A component of hull resistance is commonly modeled by Mitchell's integral, which contributes a large portion of the wave resistance experienced by a hull. Due to complexity, Mitchell's integral was excluded from the equation model and replaced by the term $\sin^2(12\text{Fr}^2)$, explaining the difference in magnitude. The data also demonstrated outliers like a drop in average total resistance when going from 5 m s-1 to 10 m s-1, decreasing from 70.986 N to 20.149 N when using the equation model. This dip can be attributed to the wave-loss term containing a squared sinusoidal wave having values ranging from 0 to 1. The reason why total resistance increases with draft and velocity can be attributed to momentum, which varies directly with velocity and mass. The draft causes total resistance to increase as the total contact area of the hull increases, making the water produce waves having large wavelengths which require a large amount of energy to produce. Increases in draft also put the boat deeper in the body of fluid, leading to more drag.

Ultimately, the methodology for this research was somewhat effective in addressing the question yet demonstrated a no extrapolatable relationship. Though both models exhibited nonlinear relationships, they demonstrated quadratic behavior that cannot be extended infinitely. One strength of the experiment was the lack of external factors such as viscosity fluctuations and dynamic sinkage inhibiting the results. Additionally, coefficients of determination ranging from 0.94 to 0.99 for all models showed a high fit for the specific data sets. Limitations included a significant 98% to 99% magnitude difference between the models due to the simplification of wave-loss and viscous terms. Improvements to the equation model could be made by incorporating Mitchell's integral, principles of Newtonian fluids, and accounting for air resistance.

MATERIALS AND METHODS

Simulation Setup

The simulation model was configured using the Prelimina Resistance Calculator. Environment variables were held constant across all trials. Density was set to 1023.6\text{ kg/m}^{3}, kinematic viscosity to 9.37 times 10^{-7} m^2, and gravity acceleration to 9.81\text{ m/s}^{2}. Waterline Z coordinates were transposed to match specific draft lengths ranging from -0.23 to -0.005 m corresponding to drafts from 0.020m to 0.245m. Speed ranges were established from a minimum speed of 5 m/s to a maximum speed of 65 m/s across 13 steps to calculate resistance in increments of five.

Equation Model Calculations

The purpose of this research is to approximate R_T using the equation model and simulation model. The equation model has two terms, the viscous resistance (R_V) and the wave-loss resistance (R_W). The R_V utilizes the frictional resistance coefficient as provided by ITTC, which has been dimensionalized.

The R_W utilizes geometry, and the Froude's derivation provided by the **Π** group. The function ($\sin^2$) was used due to the frequencies being sped up, simulating waves in a realistic manner. The component of wave-loss that is was derived from the **Π** group but the coefficient of ½ was added to simulate the loss of only a portion of the resistance from the formation of waves. The R_W was also dimensionalized. Finally, the R_W and the R_V were added to produce the R_T.

ACKNOWLEDGMENTS

We acknowledge the computational resources provided by Prelimina and reference data from the International Towing Tank Conference (ITTC) that supported this study.

\

Works Cited

1.American Physical Society. “Advancing Physics.” This Month in Physics History: Osbourne Reynolds, American Physical Society, 4 August 2024, https://www.aps.org/archives/publications/apsnews/202003/history.cfm. Accessed 8 July 2026.

2.“Geological Digressions.” Fluid flow: Froude and Reynolds numbers, Geological Digressions, 16 February 2026, https://www.geological-digressions.com/fluid-flow-froude-and-reynolds-numbers/. Accessed 20 January 2026.

3.Ghani, Aminuddin Ab. “EFFECTS OF BACKWATER ON HYDRAULIC PERFORMANCE EVALUATION OF RAINSMART MODULES IN SUSTAINABLE DRAINAGE SYSTEMS.” Research Gate, November 2018, https://www.researchgate.net/publication/329306031_EFFECTS_OF_BACKWATER_ON_HYDRAULIC_PERFORMANCE_EVALUATION_OF_RAINSMART_MODULES_IN_SUSTAINABLE_DRAINAGE_SYSTEMS. Accessed 17 June 2026.

4.IdealSimulations. “idealsimulations.” Reynold's Number, https://www.idealsimulations.com/resources/reynolds-number/. Accessed 10 July 2026.

5.International Towing Tank Conference. “ITTC – Recommended Procedures.” Testing and Extrapolation Methods High Speed Marine Vehicles Resistance Test, International Towing Tank Conference, 2002, https://ittc.info/media/2065/75-02-05-01.pdf. Accessed 13 April 2026.

6.International Towing Tank Conference. “Resistance Test.” International Towing Tank Conference, 3 September 2017, https://www.ittc.info/media/8001/75-02-02-01.pdf. Accessed 4 April 2026.

7.Nayar, Kishor G., et al. “SEAWATER THERMOPHYSICAL PROPERTIES LIBRARY.” 2017 MIT Seawater Properties Table r2, MIT, 3 December 2023, https://web.mit.edu/seawater/2017_MIT_Seawater_Property_Tables_r2b_2023c.pdf. Accessed 7 June 2026.

8.Resolved Analytics. “William Froude - Fluid Dynamicist.” Resolved Analytics, 30 April 2012, https://www.resolvedanalytics.com/the-greats-of-fluid-dynamics/william-froude-fluid-dynamics. Accessed 7 February 2026.

9.Son, Chawg Hwan. “CFD Investigation of Resistance of High-Speed Trimaran Hull Forms.” Florida Institute of Technology, July 2001, https://repository.fit.edu/cgi/viewcontent.cgi?article=2194&context=etd. Accessed 16 February 2026.

10. WordPress. “Prelimina Resistance Calculator.” Prelimina, WordPress, 10 December 2023, https://prelimina.com/resistance/. Accessed 7 June 2026.

FIGURES AND FIGURE TITLES/CAPTIONS

Figure 1. Diagram of external components of a boat. External structural features of a motorized monohull boat including draft, beam, and bow components.

Figure 1. Diagram of external components of a boat

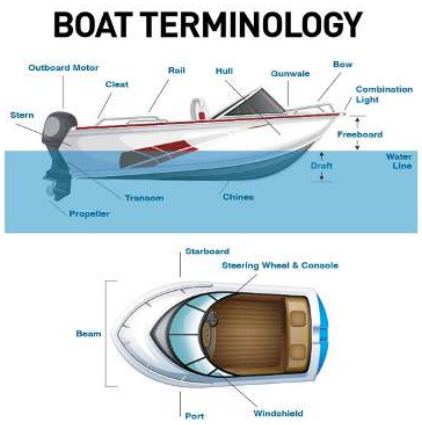


Figure 2. Table defining the different types variables mentioned.

Type of Variable	Definition	Example
Repeating Variable	- There should only be the same number of repeating variables as there are of **Π** groups. - Fundamental dimensions cannot be directly multiplied to other fundamental dimensions, but can be divided. - Each of the fundamental dimensions are expressed independently. - They cannot be problem variables nor can they be the variables that are case based; for example, g (it is specific to the Froudes number), nor can it be μ (it is specific to the Reynolds number).	When f(R_T, , μ, U, D, g) Repeating variables: , U, D
Problem Variable	- The variables that the problem is looking for with the exclusion of variables during the Buckingham Pi Theorem, though the repeating variables are included.	When f(R_T, , μ, U, D, g) Problem Variables: R_T, , μ, U, D, g
Fundamental Dimension	- The dimension in which any variable can be fundamentally written in.	Variable: R_T Variable expressed in fundamental dimensions: MLT^{-2}
Nondimensional Variable	- Variables that lack dimensions and units of any kind when expressed. - Expressed with an " * " when done through the process of nondimensionalization.	Variable: Re, u* Units: None
Dimensional Variable	- Variables defined by dimensions or units.	Variable: R_T Units: Newtons (N)
Sub-Variable	- Variables within a larger variable, usually make up the larger variable or constant that they express.	Larger Variable: Re Sub-Variable: μ, D, U
Characteristic Variable	- A variable specific to the characterized system that it is portraying, used when converting actual variables to their dimensional format used during nondimensionalization.	Actual Variable: u Characterized Variable: U

Figure 3. Table detailing all constants within the research.

Variable	**How it is Controlled**	**Justification**
The Hull Model and Type	Monohull-Wigley Hull	Has a simplistic, symmetrical shape.

Kinematic Viscosity of Sea Water	9.37×10^{-7} m^2/s	Average viscosity of seawater
Dynamic Trimming and Sinkage	None	Makes an idealized situation of the boat
Gravitational Constant on Earth	-9.81 m/s^2	Scientifically accurate
Temperature	25℃	Average temperature of seawater
Salinity	35 g/kg	Average salinity of seawater
Density	1023.6 kg/m^3	Average density of seawater sample

Figure 4. Table detailing problem variables within the research.

Numerical Variables	**Meaning**	**Equation/Units**
ρ	Water Density	kgm3
D	Boat Characteristic Draft	m
U	Boat Characteristic Velocity	ms
μ	Dynamic Viscosity	m2s
g	Gravitational Coefficient	ms2
Re	Reynold's Number	UD
Fr	Froude's Number	UgD
R_W	Wave-Resistance	$\sin^2$(12Fr2) gSD
R_V	Viscous-Resistance	0.075(log10Re-2)2 SU2
R_T	Total-Resistance	$R_V + R_W$

Figure 5. Table detailing fundamental dimensions for all problem variables.

Problem Variables	Fundamental Dimensions (MLT)
R_T	MLT2
	ML3
U	LT
D	L
μ	MLT
g	LT2

Figure 6. Image showing simulation environment section**.**
User interface input parameters for density, kinematic viscosity, and gravity acceleration within the Prelimina Resistance Calculator.

	ML3
U	LT
S	L^2
D	L
g	LT2

Figure 7. The Waterline Z Coordinate corresponding draft length.

Waterline Z Coordinate	**Draft Length**
-0.23 m	0.020 m
-0.2 m	0.050 m
-0.1 m	0.150 m
-0.05 m	0.200 m
-0.04 m	0.210 m
-0.03 m	0.220 m
-0.01 m	0.240 m
-0.005 m	0.245 m

Figure 8. Total resistance data from equation model for velocity-draft pairings.

	Hull Velocities							
Draft Lengths	**5 m s^{-1}**	**10 m s^{-1}**	**15 m s^{-1}**	**25 m s^{-1}**	**35 m s^{-1}**	**45 m s^{-1}**	**55 m s^{-1}**	**65 m s^{-1}**
0.020 m	0.2047 N	0.3046 N	0.6267 N	1.6106 N	3.0087 N	4.8022 N	6.9788 N	9.5289 N
0.050 m	2.2912 N	1.8805 N	3.7152 N	9.5073 N	17.757 N	28.341 N	41.186 N	53.26 N
0.150 m	41.305 N	15.086	26.39 N	66.601 N	124.31 N	198.39 N	288.31 N	392.65 N
0.200 m	82.192 N	24.453 N	40.506 N	101.52 N	189.44 N	302.32 N	439.32 N	599.86 N
0.210 m	91.756 N	26.389 N	43.258 N	108.27 N	202.02 N	322.39 N	468.5 N	639.69 N
0.220 m	101.46 N	28.26 N	45.849 N	114.6 N	213.81 N	341.21 N	495.84 N	677.03 N
0.240 m	121.69 N	31.951 N	50.797 N	126.63 N	236.21 N	376.96 N	547.78 N	747.96 N

Figure 9. Total resistance data from simulation model for velocity-draft pairings.

	5 m s^{-1}	**10 m s^{-1}**	**15 m s^{-1}**	**25 m s^{-1}**	**35 m s^{-1}**	**45 m s^{-1}**	**55 m s^{-1}**	**65 m s^{-1}**
0.020 m	6.163 N	22.05 N	46.64 N	120.2 N	224.6 N	358.4 N	520.8 N	711.1 N
0.050 m	14.49 N	51.57 N	109.1 N	281.3 N	525.7 N	839.2 N	1220 N	1665 N
0.150 m	74.74 N	218.6 N	455.7 N	1174 N	2193 N	3501 N	5088 N	6948 N

0.200 m	112.3 N	287.4 N	591.2 N	1521 N	2842 N	4537 N	6594 N	9004 N
0.210 m	115.4 N	290.5 N	596.3 N	1534 N	2866 N	4575 N	6650 N	9080 N
0.220 m	128.6 N	312.4 N	638.8 N	1643 N	3070 N	4900 N	7121 N	9724 N
0.240 m	142.1 N	335.9 N	684.3 N	1759 N	3287 N	5246 N	7625 N	10411 N
0.245 m	144.8 N	336.7 N	684.4 N	1759 N	3287 N	5246 N	7625 N	10411 N

Figure 10. Froudes numbers for simulation and equation model for drafts at 0.020 m.

		5 m s^{-1}	**10 m s^{-1}**	**15 m s^{-1}**	**25 m s^{-1}**	**35 m s^{-1}**	**45 m s^{-1}**	**55 m s^{-1}**	**65 m s^{-1}**
Prelimina	**0.020 m**	0.798	1.597	2.395	3.992	5.588	7.185	8.782	10.38
Equation	**0.020 m**	0.798	1.596	2.395	3.991	5.587	7.184	8.780	10.38

Figure 11. Table detailing average total resistance and standard deviation for equation and simulation model with draft length varying.

	Equation		**Prelimina**
Draft Lengths	**Average**	**Standard Deviation**	**Average**
0.020 m	3.383 N	3.447 N	251.24 N
0.050 m	19.742 N	19.444 N	588.30 N
0.150 m	114.130 N	137.833 N	2456.63 N
0.200 m	222.451 N	208.238 N	3186.11 N
0.210 m	237.784 N	221.625 N	3213.40 N
0.220 m	252.257 N	234.101 N	3442.22 N
0.240 m	279.997 N	257.643 N	3686.29 N
0.245 m	286.877 N	263.424 N	3686.74 N

Figure 12. Table detailing average total resistance and standard deviation for equation and simulation model with hull velocity varying.

	Equation		**Prelimina**
Hull Velocities	**Average**	**Standard Deviation**	**Average**
5 m s^{-1}	70.986 N	50.431 N	92.324 N
10 m s^{-1}	20.149 N	12.969 N	231.890 N
15 m s^{-1}	32.896 N	20.544 N	475.805 N
25 m s^{-1}	82.291 N	51.192 N	1223.938 N
35 m s^{-1}	153.537 N	95.496 N	2286.912 N
45 m s^{-1}	245.023 N	152.395 N	3650.325 N
55 m s^{-1}	356.064 N	221.456 N	5305.475 N
65 m s^{-1}	485.677 N	303.030 N	7244.262 N

Figure 13. Percent change from simulation model to equation model for velocity-draft pairings.

	Percentage Error (Equation - Prelimina)/ Prelimina							
Draft Lengths	**5 m s^{-1}**	**10 m s^{-1}**	**15 m s^{-1}**	**25 m s^{-1}**	**35 m s^{-1}**	**45 m s^{-1}**	**55 m s^{-1}**	**65 m s^{-1}**
0.020 m	-96.7%	-98.6%	-98.7%	-98.7%	-98.7%	-98.7%	-98.7%	-98.7%
0.050 m	-84.2%	-96.4%	-96.6%	-96.6%	-96.6%	-96.6%	-96.6%	-96.8%
0.150 m	-44.7%	-93.1%	-94.2%	-94.3%	-94.3%	-94.3%	-94.3%	-94.3%

0.200 m	-26.8%	-91.5%	-93.1%	-93.3%	-93.3%	-93.3%	-93.3%	-93.3%
0.210 m	-20.5%	-90.9%	-92.7%	-92.9%	-93.0%	-93.0%	-93.0%	-93.0%
0.220 m	-21.1%	-91.0%	-92.8%	-93.0%	-93.0%	-93.0%	-93.0%	-93.0%
0.240 m	-14.4%	-90.5%	-92.6%	-92.8%	-92.8%	-92.8%	-92.8%	-92.8%
0.245 m	-12.3%	-90.2%	-92.4%	-92.6%	-92.6%	-92.6%	-92.6%	-92.6%

Figure 14. Draft lengths vs average total resistance for equation model. Scatter plot and quadratic regression curve ($y = -22.3 + 88.8x - 6.36x^2$, $R^2 = 0.957$) showing the relationship between draft length and average total resistance calculated analytically.

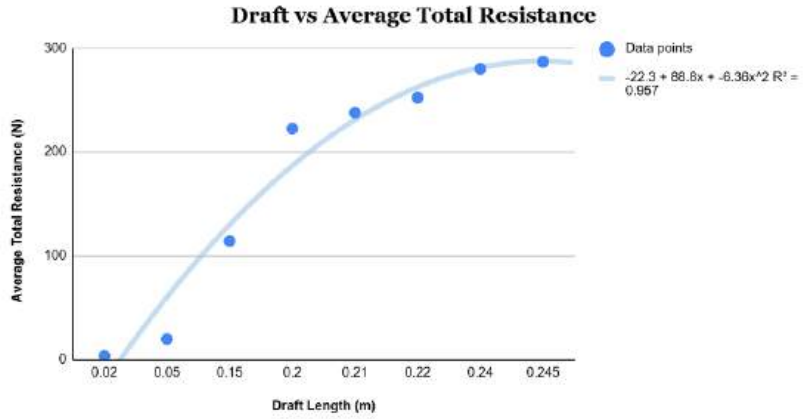


Figure 15. Hull velocities vs average total resistance for equation model. Scatter plot and quadratic regression curve ($y = 54.6 - 31.2x + 13.5x^2$, $R^2 = 0.995$) showing the relationship between hull velocity and average total resistance calculated analytically.

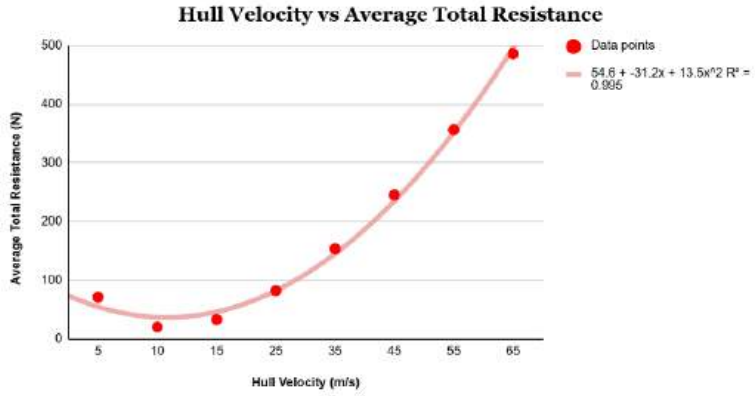


Figure 16. Draft lengths vs average total resistance for simulation model. Scatter plot and quadratic regression curve ($ y = 48.2 + 1250x - 106x^2$, $R^2 = 0.947$) showing the relationship between draft length and average total resistance derived from simulation.

Timescape Cosmology: An Alternative to the ΛCDM Standard Cosmological Model?
By Bogdan Stoiljkovic

Abstract

The Λ-Cold-Dark-Matter (ΛCDM) model is generally regarded as the standard cosmological model, although it is increasingly challenged by observations. Timescape cosmology proposes an alternative scenario, accounting for the inhomogeneity of the universe and explaining the observed acceleration of cosmic expansion without the need for Dark Energy. This paper investigates the feasibility of Timescape cosmology as an alternative to ΛCDM, particularly based on recent data which challenges the latter. While a previous investigation utilizing Hubble Flow variation appeared to favor ΛCDM, limitations with the modelling of Timescape cosmology render the results inconclusive. Furthermore, analysis of recent Supernova data from Pantheon+ provides evidence favoring Timescape instead. Based on this analysis, in this paper, I posit that Timescape cosmology could prove to be a possible alternative to ΛCDM and should receive further investigation. This paper provides a comprehensive review of current studies into Timescape cosmology. Having the potential to alleviate the need for Dark Energy, and to solve the Hubble tension, if proven preferable, Timescape cosmology would have substantial implications for future cosmological research.

Section 1: Introduction

The cosmological principle asserts that at large scales, the universe is statistically homogeneous and isotropic, with matter uniformly distributed throughout it. This principle is the basis of Friedmann–Lemaître–Robertson–Walker (FLRW) models, such as Λ-Cold-Dark-Matter (ΛCDM), which has become the standard cosmological model. While the universe is homogeneous at scales of 150-300 Mpc and beyond (Wiltshire, "Gravitational Energy as Dark Energy"; Hogg et al.; Scrimgeour et al.), at scales below this, inhomogeneity is observed in the universe – with large-scale structures of galactic filaments and walls (overdense regions) and empty voids (underdense regions) emerging, rather than uniformly distributed matter (Hoyle and Vogeley, "Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog"; Wiltshire, "Dark Energy without Dark Energy"; Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"). Furthermore, the inhomogeneity is increasing, with the universe becoming increasingly void-dominated as the gravitational pull of overdense regions attracts matter away from underdense regions, causing the latter to expand (Hoyle and Vogeley, "Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog"; Hoyle and Vogeley, "Voids in the Two‐Degree Field Galaxy Redshift Survey"). The effects of this inhomogeneity are not negligible (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Buchert et al.) although they are not sufficiently accounted for by ΛCDM. With ΛCDM facing increasing tension, particularly in light of recent data (Abdul Karim et al.; Li et al.), it may be time to search for an alternative.

A number of potential alternatives have been proposed, among them Timescape cosmology, which will remain the sole focus of this paper. Timescape cosmology offers precisely that: an alternative model of the universe which accounts for the inhomogeneity and, in so doing, alleviates the need for dark energy – one of the largest problems in modern cosmology (Wiltshire, "Gravitational Energy as Dark Energy"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). In a simplified description of Timescape cosmology, the gravitational potential is substantially higher in overdense regions than in underdense regions. Thus, as per the theory of general relativity, this would lead to time flowing more slowly in regions of overdensity and more rapidly in regions of underdensity (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"). Timescape cosmology proposes that this time dilation is the cause of the observed accelerating expansion of the universe, rather than dark energy, as is generally accepted (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"). Aside from eliminating dark energy, Timescape also claims to resolve the Hubble tension (Wiltshire, "Solution to the Cosmological Constant Problem"; Abdalla et al.). Hubble tension is one of the major tensions with ΛCDM, and refers to the discrepancy in the Hubble constant (H_0) measured through observations of the Cosmic Microwave Background (CMB) and Type Ia Supernovae (SNe Ia), with the latter yielding higher values for the rate of cosmic expansion than expected by the CMB observations (Abdalla et al.; Verde et al.; Planck Collaboration, Aghanim, et al.; Riess, Anand, et al.). Thus, if Timescape cosmology is proven to be a viable model of the universe, its implications for cosmological research would be major.

Making such bold claims, Timescape is certainly not without skepticism. Concerns exist regarding the significance of the backreactions due to inhomogeneity; while Buchert et al. (Buchert et al.) contest claims that these backreactions are irrelevant, it remains uncertain whether they are substantial enough to account for dark energy entirely. In other words, Timescape requires that time in underdense regions flows 38% faster than in overdense regions, such as in the Milky Way (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit") – it is disputed whether the gravitational potential wells of overdense regions are substantial enough to produce such high time dilation. However, a recent study by Seifert et al. (Seifert et al.), also read (Lane et al.), found that the Timescape model provided a better fit than ΛCDM to SNe Ia data from the Pantheon+ dataset. This is especially true when data from the late universe, where inhomogeneity is most present (redshift cutoff $z_{min} < 0.04$), are considered. This provides substantial support for Timescape cosmology and appears to have invigorated newer research.

This paper thus aims to investigate the viability of Timescape cosmology as a potential alternative to ΛCDM by reviewing both current and foundational literature in the field. I analyze numerous papers on ΛCDM and its limitations, as well as papers on various aspects of Timescape cosmology: the theory of the model, its limitations and opposition, and any evidence supporting the model. Based on the reviewed literature, there is insufficient evidence as of yet to disprove ΛCDM or prove Timescape, although there exists evidence supporting its validity, and as such I argue it is a field deserving of further research. Further observational tests of the

Timescape model, such as the upcoming data releases of the Euclid Space Telescope – which would provide high-resolution Baryon Acoustic Oscillation (BAO) measurements that could be used to distinguish between ΛCDM and Timescape cosmology (Wiltshire, "Solution to the Cosmological Constant Problem"; Euclid Collaboration et al.) – are necessary to further determine the viability of Timescape cosmology.

Section 2 of this paper will provide an overview of dark energy, alongside an explanation of the ΛCDM standard cosmological model and the tensions it faces. Section 3 explains Timescape cosmology and its principles in more detail, and provides evidence to support it, namely the study of Pantheon+ data by Seifert et al. (Seifert et al.). Section 4 elaborates on the concerns with Timescape cosmology, including tests performed by Saulder et al. (Saulder et al.); and Section 5 contains final discussions, including those on further observational tests for Timescape cosmology, followed by some concluding remarks.

Section 2: Dark energy, the cosmological principle, and ΛCDM

To understand Timescape cosmology, it is necessary to first understand the concept it seeks to replace – dark energy. Type Ia Supernovae (SNe Ia) have been observed to produce extremely consistent peak luminosities, leading to their use in measuring distances in the universe (Branch and Miller; Miller and Branch). However, observations by Riess et al. (Riess, Filippenko, et al.) and Perlmutter et al. (Perlmutter et al.) have found that SNe Ia at higher redshifts are significantly further than would be ordinarily expected, by 10-15% (Riess, Filippenko, et al.). This finding implied that the expansion of the universe was not constant but was, in fact, accelerating. Dark energy was eventually devised to explain this phenomenon by driving the acceleration of cosmic expansion. Further investigations by the Wilkinson Microwave Anisotropy Probe (WMAP), which analyzed the Cosmic Microwave Background (CMB), found that the universe has a flat geometry, indicating that the total mass-energy content of the universe equals the critical density (Spergel et al.). As the content of matter and dark matter in the universe comprised only 28% of the critical density, it was derived that the remaining 72% of the mass-energy of the universe must consist of dark energy (Spergel et al.). Further investigations by the Planck space telescope using similar methods refined this value to 68.3% (Planck Collaboration, Ade, et al.).

It is unknown what dark energy truly is, although there are several potential explanations (See (Ratra and Peebles; Carroll, "Quintessence and the Rest of the World"; Marcondes)). These other explanations of dark energy are outside the scope of this paper and will not be discussed. The simplest and most prominent is the cosmological constant, Λ, which Einstein added to his field equations of general relativity to maintain a static universe – one which would not expand or contract (Carroll, "The Cosmological Constant"; Peebles and Ratra). Observations of cosmic expansion (Hubble) directly contradicted the cosmological constant, resulting in it being largely abandoned from cosmological models until the aforementioned discovery of accelerating cosmic expansion. The cosmological constant, by providing a vacuum energy, results in a repulsive force capable of explaining the accelerating expansion. However, the magnitude of this vacuum energy

is under tension, with the values of the cosmological constant (supported by observation) being substantially smaller, by more than 40 orders of magnitude, than values derived by theoretical estimates, thereby creating the cosmological constant problem (Rugh and Zinkernagel). These observations of dark energy and its treatment as a cosmological constant formed the basis for the Λ-Cold-Dark-Matter (ΛCDM) model, which has become the standard cosmological model (Planck Collaboration, Aghanim, et al.).

Besides the cosmological constant, the Λ-Cold-Dark-Matter (ΛCDM) cosmological model is based on the cosmological principle, which assumes that the universe is statistically homogeneous and isotropic – that it has a constant density and is uniform when observed in all directions (Aluri et al.; Sylos Labini and Galoppo) at large scales – of 150-300 Mpc and greater (Wiltshire, "Gravitational Energy as Dark Energy"; Hogg et al.; Scrimgeour et al.). These assumptions led to the development of the Friedmann–Lemaître–Robertson–Walker (FLRW) metric, which serves as the basis for standard cosmological models such as ΛCDM (Planck Collaboration, Aghanim, et al.; Aluri et al.). However, the validity of the cosmological principle, and in particular the suitability of ΛCDM in modeling the universe, are being increasingly challenged by numerous means (Aluri et al.; Sylos Labini and Galoppo). Most prominently in regard to Timescape cosmology is the apparent inhomogeneity of the universe. The universe is, in fact, dominated by cosmic voids (Hoyle and Vogeley, "Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog"; Hoyle and Vogeley, "Voids in the Two‐Degree Field Galaxy Redshift Survey"; Sylos Labini and Galoppo), surrounded by large-scale structures such as cosmic filaments and walls – including the CfA2 Great Wall (Geller and Huchra) and the Sloan Great Wall (Gott et al.). Large-scale sky surveys such as the Sloan Digital Sky Survey (SDSS) (Gott et al.) and the more recent Dark Energy Spectroscopic Instrument (DESI) survey (DESI Collaboration et al.) have been particularly useful in revealing this, as shown in Figure 1. The presence and prominence of such large-scale structure in the nearby universe may have more significant effects in regard to what is observed than assumed by ΛCDM and other FLRW-based models limited by the cosmological principle (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Sylos Labini and Galoppo; Ginat and Ferreira).

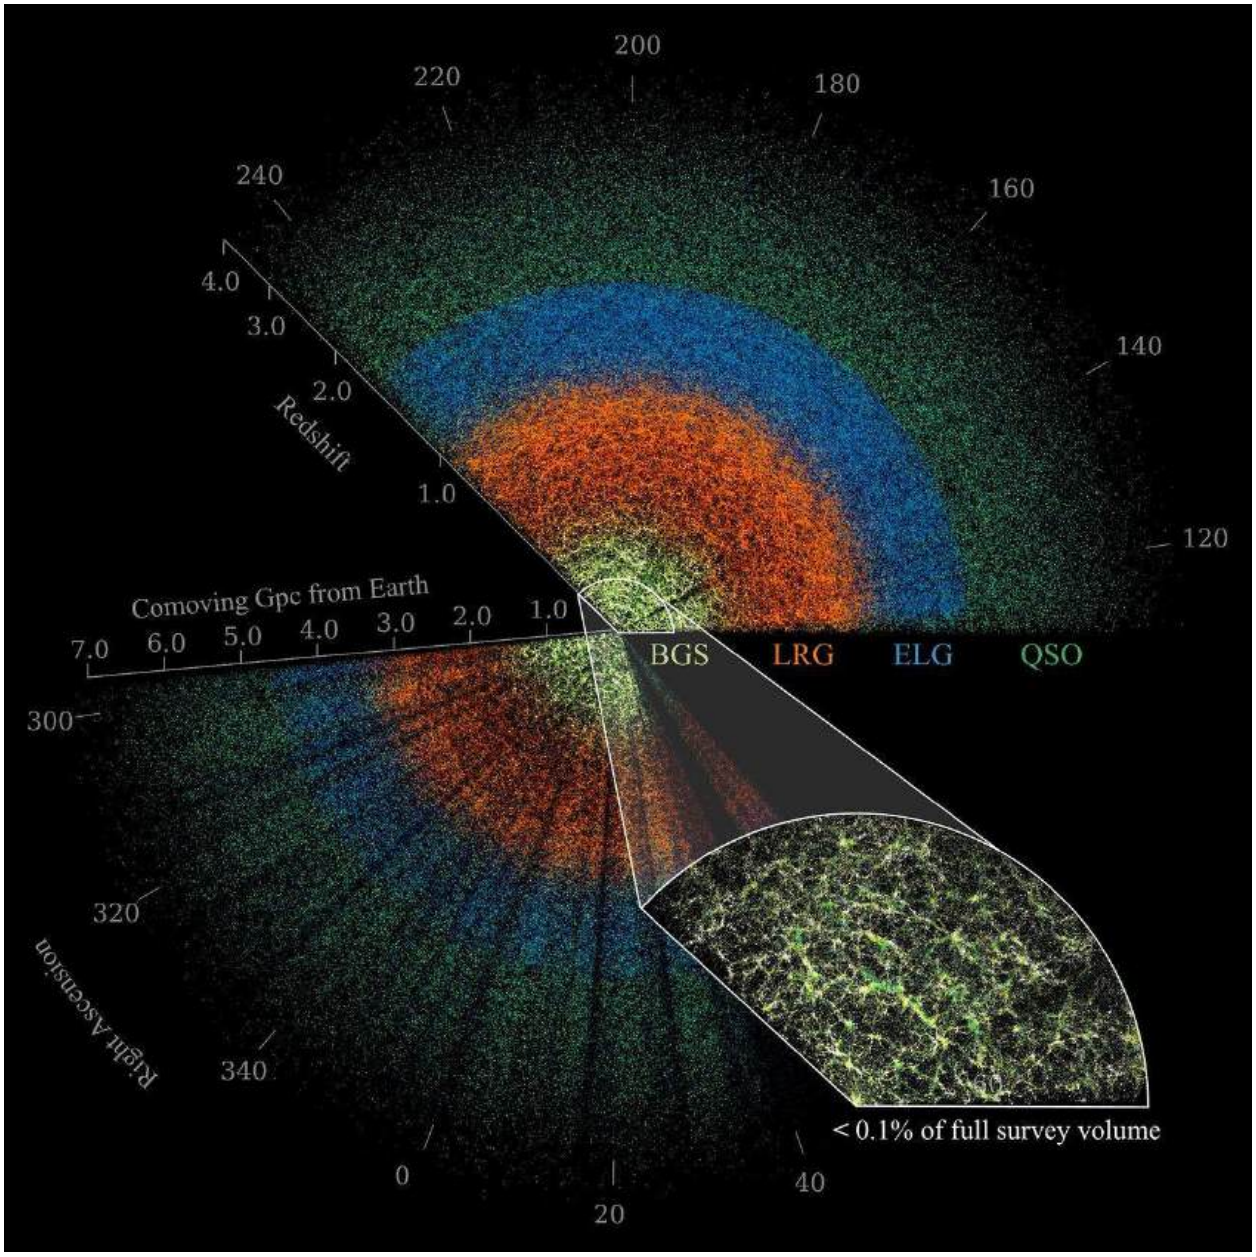


Fig. 1: Courtesy of NOIRLab (DESI Collaboration/DOE/KPNO/NOIRLab/NSF/AURA/C. Lamman). A map of the universe produced by the Dark Energy Spectroscopic Instrument (DESI). Each point represents one of various objects: yellow represents Bright Galaxies (BGS), orange represents Luminous Red Galaxies (LRG), blue represents Emission-Line Galaxies (ELG), and green represents Quasars (QSO). While the universe appears homogeneous at high redshifts, clear large-scale structure and inhomogeneity are present in the nearby universe, as shown by the inset image.

Further tensions with ΛCDM from the DESI survey arise in the form of the support the latter has found for evolving dark energy (Abdul Karim et al.; Li et al.). Evidence for evolving dark energy is in direct opposition to the cosmological constant, which assumes dark energy to be a constant vacuum energy that accelerates the expansion of the universe at a constant rate, and

thus also strongly challenges the ΛCDM model. These results instead find that a $w_0 w_a CDM$ model is preferred over ΛCDM (Abdul Karim et al.; Mena-Fernández et al.). Yet another major tension with ΛCDM is the Hubble tension: a discrepancy in the measured Hubble constant (H_0). Observations of Cepheid variables – regularly-pulsating stars in the nearby 'late' universe – and Type Ia Supernovae (SNe Ia) by the Hubble Space Telescope (HST) and James Webb Space Telescope (JWST) both provide a Hubble constant of $H_0 = 73.0 \pm 1.0 km\, s^{-1} Mpc^{-1}$ (Riess, Anand, et al.), while observations of the Cosmic Microwave Background (CMB) in combination with ΛCDM cosmology provide a Hubble constant of $H_0 = 67.4 \pm 0.5 km\, s^{-1} Mpc^{-1}$ (Planck Collaboration, Aghanim, et al.). This is a consistent 5σ tension that seems to arise from the ΛCDM cosmological model, which, in combination with the other aforementioned tensions, suggests that a more suitable cosmological model may be necessary (Di Valentino et al.; Abdalla et al.).

Section 3: An overview of Timescape cosmology

The inhomogeneity of the universe outlined in the previous section can have quite substantial effects. Namely, overdense regions of galactic walls and filaments contain more matter (baryonic and dark) than underdense regions, and as per general relativity, this difference in matter density results in differential curvature of spacetime (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit"). More specifically, overdense regions will result in greater spacetime curvature, thereby resulting in time flowing more slowly there, while underdense regions will have a negative curvature of spacetime, resulting in time there flowing more rapidly (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit"). This effect is fundamental to the backreaction of inhomogeneity and is integral to the Timescape model, which argues that dark energy is instead a misidentification of these effects, and, as such, the observations of accelerating cosmic expansion (Riess, Filippenko, et al.; Perlmutter et al.) are apparent rather than fundamental (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). Timescape cosmology claims that the different time flow of overdense and underdense regions leads to differential expansion of the two – underdense regions, with their time flowing more quickly, will be older than overdense regions and will thus have expanded more (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). This means that light traveling through underdense regions will be redshifted more than light passing through overdense regions, as is shown by Figure 2.

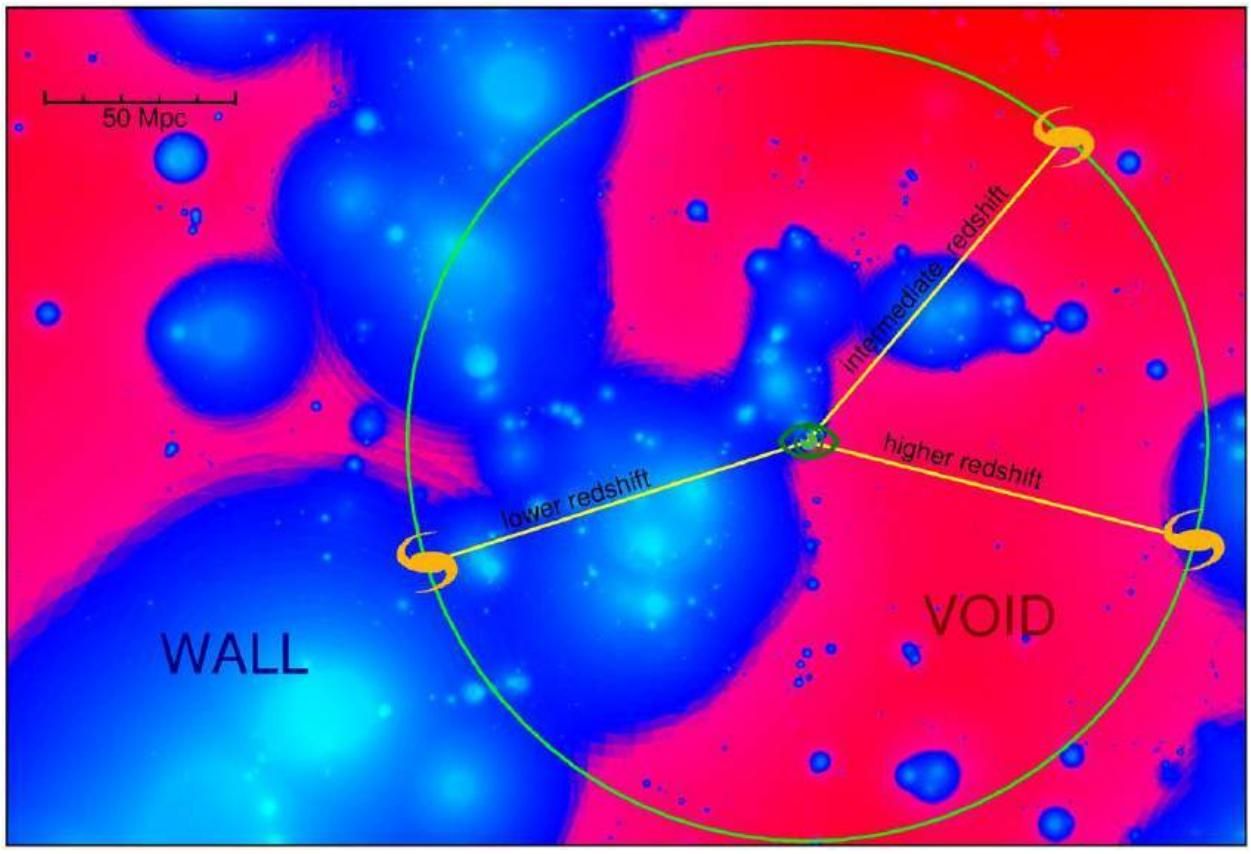


Fig. 2: Extracted from Saulder et al. (Saulder et al.) Figure 1. A diagram showing the different redshifts observed from galaxies (represented by the yellow galaxy symbols) at the same distance from the observer (represented by the green eye symbol) due to the inhomogeneity of the universe as predicted by Timescape. Galaxies separated from the observer by underdense regions (voids, represented by red) will have a higher observed redshift than galaxies separated from the observer by overdense regions (galactic filaments and walls, represented by blue).

This principle allows Timescape cosmology to have some significant implications. Gravitational attraction results in overdense regions attracting matter from underdense regions, causing the underdense regions to grow and thereby resulting in an increasing void fraction in the late universe (Hoyle and Vogeley, "Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog"; Hoyle and Vogeley, "Voids in the Two‐Degree Field Galaxy Redshift Survey"). As explained above, Timescape cosmology posits that light passing through underdense regions will, in effect, have a higher redshift than light passing through exclusively overdense regions. Thus, the increasing void fraction in the late universe would correspond to higher redshift observations of distant objects, as observed by Reiss et al. (Riess, Filippenko, et al.) and Perlmutter et al. (Perlmutter et al.), thereby explaining the apparent acceleration of cosmic expansion without the need for dark energy. Furthermore, Timescape posits that spherical averaging of this higher redshift due to the faster time flow in underdense regions compared to that of overdense regions (in which the Milky Way resides, and which is thus our reference point) results in the Hubble constant (H_0) measured by local observations being higher than that inferred by the Cosmic Microwave Background (CMB) (Wiltshire, "Solution to the

Cosmological Constant Problem"). This would imply that Timescape cosmology, by accounting for this, resolves the Hubble tension (Wiltshire, "Solution to the Cosmological Constant Problem"). Finally, of note is the observation that acceleration of cosmic expansion begins at a period (redshift) similar to the emergence of the large-scale structure utilized by Timescape cosmology (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). The implications of this alone appear to be in support of Timescape cosmology. As stated previously, by alleviating the need for dark energy, which remains one of the biggest problems in modern cosmology, and resolving the Hubble tension, Timescape cosmology would greatly affect all future cosmological studies if it is found to be accurate.

There exists a substantial lack of research into the viability of Timescape cosmology through observational tests. One such study is the recent analysis of Type Ia Supernova (SNe Ia) data from the Pantheon+ database (Scolnic et al.) by Seifert et al. (Seifert et al.); also see (Lane et al.). The analysis computes a distance modulus for the SNe Ia from the Pantheon+ observational data and compares it with a corresponding theoretical distance modulus produced by both ΛCDM and Timescape cosmology. Extreme care was taken to ensure that the observational distance moduli were derived independently of assumptions from cosmological models to ensure an effective comparison (Seifert et al.; Lane et al.). To do so, the peculiar velocity correction and other corrections (such as that for Malmquist bias) were omitted from the analysis, having been deemed insignificant in regard to the Bayes factor produced and impossible to include while preserving an effective comparison of the two models (Seifert et al.). Following this, the Bayes factor for the two models was produced, shown by Figure 3, finding preference for Timescape cosmology over ΛCDM (Seifert et al.). More specifically, the data is separated into a number of different redshift cutoffs (z_{min}). Higher preference is found at lower redshift cutoffs ($z_{min} \lesssim 0.02$) where the data contains SNe Ia at lower redshifts that are more significantly affected by inhomogeneity; while higher redshift cutoffs ($z_{min} \gtrsim 0.04$) contain only the more distant SNe Ia which are less affected by the inhomogeneity, and thus a lower preference for Timescape cosmology is observed (Seifert et al.). Interestingly, Timescape cosmology is still slightly preferred over ΛCDM even at higher redshift cutoffs ($z_{min} \gtrsim 0.075$), suggesting it is capable of effectively modeling the expansion of the universe even at these higher scales and thereby further supporting the validity of Timescape cosmology (Seifert et al.). This analysis provides support for Timescape cosmology, showing that it may be an effective alternative to ΛCDM, although it does not alone prove Timescape cosmology or disprove ΛCDM. Further investigations akin to this one, particularly testing Timescape using other observables, would be necessary, although this analysis nevertheless shows that Timescape cosmology shows promise not just in theory but also in observational data.

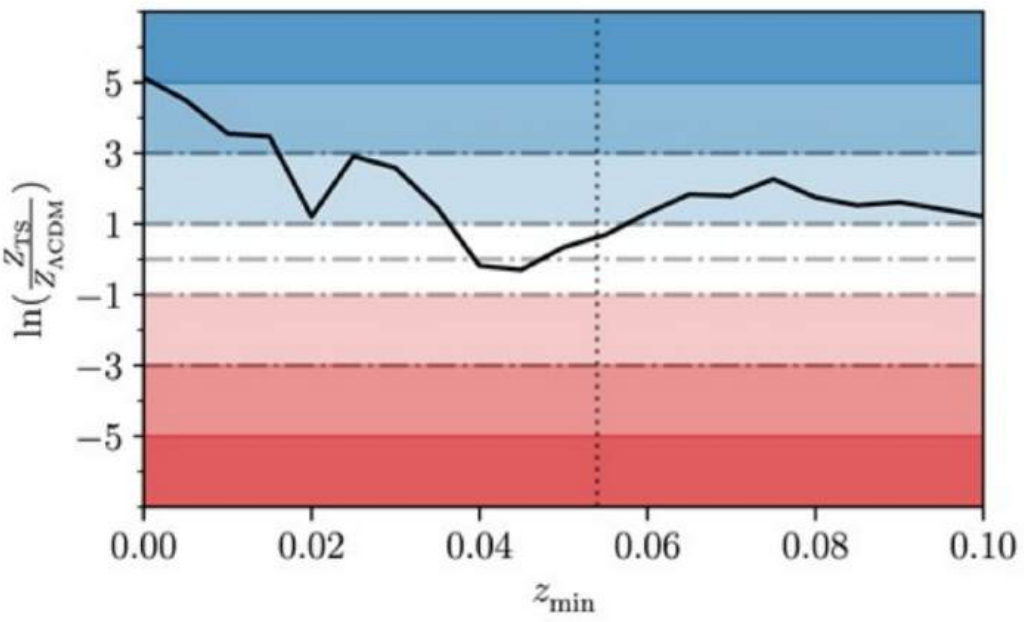


Fig. 3: Adapted from Seifert et al. (Seifert et al.) Figure 1. A Bayes factor plot comparing the preference for Timescape (shown by the blue shaded region, $ln(\frac{z_{TS}}{z_{\Lambda CDM}}) > 1$) and the preference for ΛCDM (shown by the red shaded region, $ln(\frac{z_{TS}}{z_{\Lambda CDM}}) < -1$). The unshaded region favors neither model. As can be seen, there is substantial preference for Timescape cosmology over ΛCDM at low redshift cutoffs (z_{min}), with this preference decreasing as the redshift cutoff increases and the closer supernovae – which are more strongly affected by the inhomogeneity – are excluded from the analysis, reaching a minimum at $z_{min}\sim 0.045$. Interestingly, the preference for Timescape increases again as the redshift cutoff is increased further, with a local maximum at $z_{min}\sim 0.075$. If distances of 150-300 Mpc and beyond – equivalent to a redshift cutoff of $z_{min}\sim 0.035$ to $z_{min}\sim 0.070$ – are considered to be the scale at which the universe is homogeneous (Wiltshire, "Gravitational Energy as Dark Energy"; Hogg et al.; Scrimgeour et al.), the preference for Timescape cosmology at and above this redshift cutoff appears to be in contradiction with the cosmological principle.

Section 4: Concerns with Timescape cosmology

Being such a large paradigm shift, Timescape cosmology is not without concern and opposition. As explained previously, there is presently a distinct lack of research into Timescape cosmology, meaning that reported opposition to it is also limited. One major concern with Timescape cosmology is the significance of the backreaction of the inhomogeneities central to it. The significance of these backreactions has been a matter of much debate, and a number of previous papers (Ishibashi and Wald; Green and Wald) have claimed them to be negligible. This has been contested by Buchert et al. (Buchert et al.), and the results of Seifert et al. (Seifert et al.) further oppose these claims, by demonstrating that accounting for inhomogeneities produces a preferable fit to observational data than ignoring them, thereby demonstrating that the backreactions of inhomogeneity have a measurable effect and are not negligible. However, the question of whether the backreactions are significant enough to fully account for dark energy as

proposed by Timescape cosmology still remains. Timescape cosmology requires that the time dilation caused by the difference in density between regions of overdensity and underdensity results in time flowing 38% more rapidly in the latter (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit") – an extremely high value. Furthermore, this would result in the age of the universe varying by billions of years between the different regions (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"), the viability of which is similarly debated.

A second major concern with Timescape cosmology arises from a cosmological test by Saulder et al. (Saulder et al.) utilizing Hubble Flow Variations as a test to determine the viability of Timescape cosmology. This was accomplished by directly comparing Timescape cosmology to ΛCDM with respect to the expansion rates of voids and galactic walls. To do so, the Hubble parameters of low-redshift ($z<0.1$) galaxies were collected from the Sloan Digital Sky Survey (SDSS) and 2MASS Redshift Survey (2MRS) datasets, by comparing the redshifts of galaxies with their fundamental planes. These observations were then compared to simulated data from the Millennium simulation using a ΛCDM model, and an approximation of Timescape cosmology obtained by deforming the ΛCDM model to fit bounds suggested by Timescape cosmology (Saulder et al.). This test found that the data favored the ΛCDM model, with the approximated Timescape model producing significantly worse fits to the data (Saulder et al.). However, the validity of this test was challenged by Wiltshire (Wiltshire, "Comment on 'Hubble Flow Variations as a Test for Inhomogeneous Cosmology'"), who argued that the approximated Timescape model used in the test suffered from systematic errors by failing to apply a restriction on the quasilocal Hubble flow required by Timescape cosmology, thus rendering the approximation ineffective at modeling Timescape cosmology. As such, the results of the test by Saulder et al. (Saulder et al.) are inconclusive with regard to the viability of Timescape cosmology. Furthermore, the findings of Seifert et al. (Seifert et al.) are also in direct opposition to these findings by Saulder et al. (Saulder et al.), further challenging the conclusions of the study. Both of these concerns demonstrate the distinct lack of research into Timescape cosmology to date, and the necessity for further research to definitively conclude not just whether Timescape cosmology or ΛCDM is preferable, but whether Timescape cosmology is truly a viable cosmological model or if a different alternative to ΛCDM should be pursued instead.

Conclusion and Discussions

Dark energy has remained one of the greatest cosmological mysteries since its conception, being treated simply as a cosmological constant and forming the basis of our most effective standard cosmological model, the Λ-Cold-Dark-Matter (ΛCDM) model, although remaining unknown in nature. Timescape cosmology, while a major paradigm shift, has the potential to change all this, providing an alternative explanation to the observations which led to the conception of dark energy without necessitating it (Wiltshire, "Gravitational Energy as Dark Energy"; Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire,

"Cosmic Clocks, Cosmic Variance and Cosmic Averages"). This is accomplished by accounting for the backreactions of the inhomogeneities present at certain scales of the universe (<150Mpc, (Wiltshire, "Gravitational Energy as Dark Energy"; Aluri et al.)) where large-scale structure such as galactic filaments, walls, and voids arise. The difference in the densities of these regions may have a larger effect than is accounted for by models based on the cosmological principle, which assume a statistically homogeneous and isotropic universe at large scales (Aluri et al.). By accounting for this difference in density and thus the gravitational potential, general relativity is used to determine different rates of time flow throughout the universe (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit"). This differential time flow leads to a differential age of the universe, with voids – having faster-flowing time – being considerably older than the galactic filaments and walls, and thus being more expanded (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). Light passing through this differential expansion will be redshifted by different amounts, with light passing through voids being more redshifted. This, in combination with the increasing void dominance of the late universe (Hoyle and Vogeley, "Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog"; Hoyle and Vogeley, "Voids in the Two‐Degree Field Galaxy Redshift Survey"), has been utilized by Timescape cosmology to explain the observations (Riess, Filippenko, et al.; Perlmutter et al.) that led to the apparent acceleration of cosmic expansion (Wiltshire, "Gravitational Energy as Dark Energy"; Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution"; Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"). This is also utilized to explain the Hubble tension of ΛCDM (Wiltshire, "Solution to the Cosmological Constant Problem").

There is skepticism and opposition to Timescape cosmology, particularly regarding the significance of the backreactions of inhomogeneities. Previous papers (Ishibashi and Wald; Green and Wald) arguing that the effects of the backreactions were insignificant were challenged by Buchert et al. (Buchert et al.). However, the requirements for such substantial time dilation (with time flowing 38% faster in underdense regions than in overdense ones (Wiltshire, "Cosmic Clocks, Cosmic Variance and Cosmic Averages"; Wiltshire, "Cosmological Equivalence Principle and the Weak-Field Limit")) and the substantial difference in the age of the universe (with the universe being several billion years older in regions of underdensity (Wiltshire, "From Time to Timescape - Einstein's Unfinished Revolution")) remain points of concern. A test of Timescape cosmology by Saulder et al. (Saulder et al.) compared Timescape cosmology with ΛCDM using a comparison of observed Hubble parameters of galaxies from the Sloan Digital Sky Survey (SDSS) and 2MASS Redshift Survey (2MRS) datasets to simulated data from the Millennium simulation using ΛCDM and an approximated Timescape model. The test found a preference for ΛCDM over Timescape cosmology (Saulder et al.), although systematic problems with the approximation of Timescape cosmology limited the conclusions that could be drawn from this result (Wiltshire, "Comment on 'Hubble Flow Variations as a Test for Inhomogeneous

Cosmology'"). A recent analysis of Type Ia Supernova (SNe Ia) data from the Pantheon+ dataset by Seifert et al. (Seifert et al.), however, found a preference for Timescape cosmology over ΛCDM, particularly in data including low-redshift SNe Ia which are more strongly affected by the inhomogeneity. This provides substantial support for the theory as it is in direct opposition with the results of the test by Saulder et al. (Saulder et al.) and provides evidence that backreactions of inhomogeneity are not insignificant (Ishibashi and Wald; Green and Wald). However, it alone is insufficient to prove Timescape cosmology, and there is currently a lack of research regarding it. As such, while the findings of Seifert et al. (Seifert et al.) provide support for Timescape cosmology, more research is necessary to determine if Timescape cosmology is viable, or if one of the alternative models is more suitable.

The Euclid Space Telescope (Euclid Collaboration et al.) is expected to be able to further determine the viability of Timescape cosmology – particularly by using the Clarkson-Bassett-Lu test of the Copernican Principle (Clarkson et al.) to determine if the Friedmann equation is true and FLRW models are effective, or if the predictions of Timescape are preferred (Wiltshire, "Solution to the Cosmological Constant Problem"). In the meantime, I posit that Timescape cosmology should receive further research to determine its viability, given the major implications for cosmological studies it would have if proven true. The status of ΛCDM as the standard cosmological model can be attributed to its great success in effectively modeling the universe in most regards, and as such further tests akin to Seifert et al. (Seifert et al.) and Lane et al. (Lane et al.) with other observables, such as the CMB or BAO, are necessary to develop Timescape cosmology into a concordant model or disprove it entirely. Timescape cosmology is a major paradigm shift and, as such, receives warranted skepticism. However, I argue that this should not result in it being dismissed outright, in light of arguments regarding the conventionalism of ΛCDM (Merritt). The paradigm that extraordinary claims require extraordinary evidence (Saulder et al.) may be misguided; if Timescape cosmology is able to achieve concordance with other observables equal to that of ΛCDM, while this alone would be insufficient to prove preference for Timescape cosmology, it should be sufficient to solidify Timescape as a viable alternative. Finally, the w_0w_aCDM model described by Karim et al. (Abdul Karim et al.) has shown good promise (Abdul Karim et al.; Li et al.), and may serve as a more conventional competitor to Timescape cosmology. Performing a test – again, akin to that of Seifert et al. (Seifert et al.) and Lane et al. (Lane et al.), at least in nature – would be of benefit, potentially having the ability to determine if ΛCDM should be replaced by w_0w_aCDM instead of Timescape.

Acknowledgements

I would like to thank Dr. Ryan Farber and Ms. Catherine Peretti for their help in producing this paper, providing invaluable feedback and guidance throughout the research process. I would also like to extend my appreciation to the Indigo Research team for allowing me to be a part of the Indigo Research Summer Intensive (IRIS) Program, without which I would not have produced this paper.

Works Cited

Abdalla, Elcio, et al. “Cosmology Intertwined: A Review of the Particle Physics, Astrophysics, and Cosmology Associated with the Cosmological Tensions and Anomalies.” Journal of High Energy Astrophysics, vol. 34, June 2022, pp. 49–211. arXiv.org, https://doi.org/10.1016/j.jheap.2022.04.002.

Abdul Karim, M., et al. “DESI DR2 Results. II. Measurements of Baryon Acoustic Oscillations and Cosmological Constraints.” Physical Review D, vol. 112, no. 8, Oct. 2025, p. 083515. DOI.org (Crossref), https://doi.org/10.1103/tr6y-kpc6.

Aluri, Pavan Kumar, et al. “Is the Observable Universe Consistent with the Cosmological Principle?” Classical and Quantum Gravity, vol. 40, no. 9, May 2023, p. 094001. arXiv.org, https://doi.org/10.1088/1361-6382/acbefc.

Branch, David, and Douglas L. Miller. “Type IA Supernovae as Standard Candles.” The Astrophysical Journal Letters, vol. 405, Mar. 1993, p. L5. NASA ADS, https://doi.org/10.1086/186752.

Buchert, T., et al. “Is There Proof That Backreaction of Inhomogeneities Is Irrelevant in Cosmology?” Classical and Quantum Gravity, vol. 32, no. 21, Nov. 2015, p. 215021. arXiv.org, https://doi.org/10.1088/0264-9381/32/21/215021.

Carroll, Sean M. “Quintessence and the Rest of the World.” Physical Review Letters, vol. 81, no. 15, Oct. 1998, pp. 3067–70. arXiv.org, https://doi.org/10.1103/PhysRevLett.81.3067.

———. “The Cosmological Constant.” Living Reviews in Relativity, vol. 4, no. 1, Dec. 2001, p. 1. arXiv.org, https://doi.org/10.12942/lrr-2001-1.

Clarkson, Chris, et al. “A General Test of the Copernican Principle.” Physical Review Letters, vol. 101, no. 1, July 2008, p. 011301. arXiv.org, https://doi.org/10.1103/PhysRevLett.101.011301.

DESI Collaboration, et al. “The DESI Experiment Part I: Science,Targeting, and Survey Design.” arXiv:1611.00036, arXiv, 13 Dec. 2016. arXiv.org, https://doi.org/10.48550/arXiv.1611.00036.

Di Valentino, Eleonora, et al. “In the Realm of the Hubble Tension—a Review of Solutions*.” Classical and Quantum Gravity, vol. 38, no. 15, July 2021, p. 153001. DOI.org (Crossref), https://doi.org/10.1088/1361-6382/ac086d.

Euclid Collaboration, et al. “Euclid: I. Overview of the Euclid Mission.” Astronomy & Astrophysics, vol. 697, May 2025, p. A1. DOI.org (Crossref), https://doi.org/10.1051/0004-6361/202450810.

Geller, Margaret J., and John P. Huchra. “Mapping the Universe.” Science, vol. 246, no. 4932, Nov. 1989, pp. 897–903. science.org (Atypon), https://doi.org/10.1126/science.246.4932.897.

Ginat, Yonadav Barry, and Pedro G. Ferreira. “Apparent Dark-Energy Evolution from Cosmic Inhomogeneities.” arXiv:2601.20633, arXiv, 28 Jan. 2026. arXiv.org, https://doi.org/10.48550/arXiv.2601.20633.

Gott, J. Richard, et al. “A Map of the Universe.” The Astrophysical Journal, vol. 624, no. 2, May 2005, pp. 463–84. arXiv.org, https://doi.org/10.1086/428890.

Green, Stephen R., and Robert M. Wald. “How Well Is Our Universe Described by an FLRW Model?” arXiv.Org, 30 July 2014, https://doi.org/10.1088/0264-9381/31/23/234003.

Hogg, David W., et al. “Cosmic Homogeneity Demonstrated with Luminous Red Galaxies.” The Astrophysical Journal, vol. 624, no. 1, May 2005, pp. 54–58. arXiv.org, https://doi.org/10.1086/429084.

Hoyle, Fiona, and Michael S. Vogeley. “Voids in the Point Source Catalogue Survey and the Updated Zwicky Catalog.” The Astrophysical Journal, vol. 566, no. 2, Feb. 2002, pp. 641–51. DOI.org (Crossref), https://doi.org/10.1086/338340.

———. “Voids in the Two‐Degree Field Galaxy Redshift Survey.” The Astrophysical Journal, vol. 607, no. 2, June 2004, pp. 751–64. DOI.org (Crossref), https://doi.org/10.1086/386279.

Hubble, Edwin. “A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae.” Proceedings of the National Academy of Science, vol. 15, Mar. 1929, pp. 168–73. NASA ADS, https://doi.org/10.1073/pnas.15.3.168.

Ishibashi, Akihiro, and Robert M. Wald. “Can the Acceleration of Our Universe Be Explained by the Effects of Inhomogeneities?” Classical and Quantum Gravity, vol. 23, no. 1, Jan. 2006, pp. 235–50. arXiv.org, https://doi.org/10.1088/0264-9381/23/1/012.

Lane, Zachary G., et al. “Cosmological Foundations Revisited with Pantheon+.” Monthly Notices of the Royal Astronomical Society, vol. 536, no. 2, Jan. 2025, pp. 1752–77. Silverchair, https://doi.org/10.1093/mnras/stae2437.

Li, Tian-Nuo, et al. “Dark Energy in the DESI Era: A Brief Review of Evidence, Beyond-ΛCDM Interpretations, and Tensions.” Research in Astronomy and Astrophysics, vol. 26, no. 8, July 2026, p. 084002. Institute of Physics, https://doi.org/10.1088/1674-4527/ae8429.

Marcondes, Rafael J. F. “Interacting Dark Energy Models in Cosmology and Large-Scale Structure Observational Tests.” arXiv:1610.01272, arXiv, 5 Oct. 2016. arXiv.org, https://doi.org/10.48550/arXiv.1610.01272.

Mena-Fernández, J., et al. “Dark Energy Survey: DESI-Independent Angular BAO Measurement.” Physical Review D, vol. 113, no. 12, June 2026, p. 123521. arXiv.org, https://doi.org/10.1103/d474-zzzf.

Merritt, David. “Cosmology and Convention.” Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, vol. 57, Feb. 2017, pp. 41–52. DOI.org (Crossref), https://doi.org/10.1016/j.shpsb.2016.12.002.

Miller, Douglas L., and David Branch. “Supernova Absolute-Magnitude Distributions.” The Astronomical Journal, vol. 100, Aug. 1990, p. 530. DOI.org (Crossref), https://doi.org/10.1086/115534.

Peebles, P. J. E., and Bharat Ratra. “The Cosmological Constant and Dark Energy.” Reviews of Modern Physics, vol. 75, no. 2, Apr. 2003, pp. 559–606. arXiv.org, https://doi.org/10.1103/RevModPhys.75.559.

Perlmutter, S., et al. “Measurements of Omega and Lambda from 42 High-Redshift Supernovae.” The Astrophysical Journal, vol. 517, no. 2, June 1999, pp. 565–86. arXiv.org, https://doi.org/10.1086/307221.

Planck Collaboration, P. A. R. Ade, et al. “Planck 2013 Results. I. Overview of Products and Scientific Results.” Astronomy & Astrophysics, vol. 571, Nov. 2014, p. A1. arXiv.org, https://doi.org/10.1051/0004-6361/201321529.

Planck Collaboration, N. Aghanim, et al. “Planck 2018 Results. VI. Cosmological Parameters.” Astronomy & Astrophysics, vol. 641, Sept. 2020, p. A6. arXiv.org, https://doi.org/10.1051/0004-6361/201833910.

Ratra, Bharat, and P. J. E. Peebles. “Cosmological Consequences of a Rolling Homogeneous Scalar Field.” Physical Review D, vol. 37, no. 12, June 1988, pp. 3406–27. APS, https://doi.org/10.1103/PhysRevD.37.3406.

Riess, Adam G., Gagandeep S. Anand, et al. “Crowded No More: The Accuracy of the Hubble Constant Tested with High Resolution Observations of Cepheids by JWST.” arXiv:2307.15806, arXiv, 28 July 2023. arXiv.org, https://doi.org/10.48550/arXiv.2307.15806.

Riess, Adam G., Alexei V. Filippenko, et al. “Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant.” The Astronomical Journal, vol. 116, no. 3, Sept. 1998, pp. 1009–38. DOI.org (Crossref), https://doi.org/10.1086/300499.

Rugh, Svend Erik, and Henrik Zinkernagel. “The Quantum Vacuum and the Cosmological Constant Problem.” arXiv:hep-th/0012253, arXiv, 28 Dec. 2000. arXiv.org, https://doi.org/10.48550/arXiv.hep-th/0012253.

Saulder, Christoph, et al. “Hubble Flow Variations as a Test for Inhomogeneous Cosmology.” Astronomy & Astrophysics, vol. 622, Feb. 2019, p. A83. DOI.org (Crossref), https://doi.org/10.1051/0004-6361/201629174.

Scolnic, Dan, et al. “The Pantheon+ Analysis: The Full Data Set and Light-Curve Release.” The Astrophysical Journal, vol. 938, no. 2, Oct. 2022, p. 113. DOI.org (Crossref), https://doi.org/10.3847/1538-4357/ac8b7a.

Scrimgeour, Morag, et al. “The WiggleZ Dark Energy Survey: The Transition to Large-Scale Cosmic Homogeneity.” Monthly Notices of the Royal Astronomical Society, vol. 425, no. 1, Sept. 2012, pp. 116–34. arXiv.org, https://doi.org/10.1111/j.1365-2966.2012.21402.x.

Seifert, Antonia, et al. “Supernovae Evidence for Foundational Change to Cosmological Models.” Monthly Notices of the Royal Astronomical Society: Letters, vol. 537, no. 1, Feb. 2025, pp. L55–60. Silverchair, https://doi.org/10.1093/mnrasl/slae112.

Spergel, D. N., et al. “Wilkinson Microwave Anisotropy Probe (WMAP) Three Year Results: Implications for Cosmology.” The Astrophysical Journal Supplement Series, vol. 170, no. 2, June 2007, pp. 377–408. arXiv.org, https://doi.org/10.1086/513700.

Sylos Labini, Francesco, and Marco Galoppo. “Detection of Anisotropic Cosmic Structures on a Gigaparsec Scale.” Nature, vol. 655, no. 8122, July 2026, pp. 327–31. www.nature.com, https://doi.org/10.1038/s41586-026-10702-5.

Verde, Licia, et al. “A Tale of Many H0.” Annual Review of Astronomy and Astrophysics, vol. 62, no. Volume 62, 2024, Sept. 2024, pp. 287–331. www.annualreviews.org, https://doi.org/10.1146/annurev-astro-052622-033813.

Wiltshire, David L. “Comment on ‘Hubble Flow Variations as a Test for Inhomogeneous Cosmology.’” Astronomy & Astrophysics, vol. 624, Apr. 2019, p. A12. www.aanda.org, https://doi.org/10.1051/0004-6361/201834833.

———. “Cosmic Clocks, Cosmic Variance and Cosmic Averages.” New Journal of Physics, vol. 9, no. 10, Oct. 2007, pp. 377–377. DOI.org (Crossref), https://doi.org/10.1088/1367-2630/9/10/377.

———. “Cosmological Equivalence Principle and the Weak-Field Limit.” Physical Review D, vol. 78, no. 8, Oct. 2008, p. 084032. arXiv.org, https://doi.org/10.1103/PhysRevD.78.084032.

———. “Dark Energy without Dark Energy.” Dark Matter in Astroparticle and Particle Physics, 2008, pp. 565–96. arXiv.org, https://doi.org/10.1142/9789812814357_0053.

———. “From Time to Timescape - Einstein’s Unfinished Revolution.” International Journal of Modern Physics D, vol. 18, no. 14, Dec. 2009, pp. 2121–34. arXiv.org, https://doi.org/10.1142/S0218271809016193.

———. “Gravitational Energy as Dark Energy: Cosmic Structure and Apparent Acceleration.” arXiv:1102.2045, arXiv, 10 Feb. 2011. arXiv.org, https://doi.org/10.48550/arXiv.1102.2045.

———. “Solution to the Cosmological Constant Problem.” arXiv:2404.02129, arXiv, 2 Apr. 2024. arXiv.org, https://doi.org/10.48550/arXiv.2404.02129.

A Review Analysis of the Feasibility of Implementing RGBW Sensors in Cameras

By Kiaan Shah

Abstract

RGB (Red-green-blue) Bayer colour filter arrays have been normalized for digital image sensors. However, their single-colour sampling method and narrow spectral range hinder image quality in low-light conditions and cause undesired side effects. RGBW sensors attempt to fix this issue by replacing a green pixel in each array with a white pixel, which is meant to facilitate greater light transmission as well as improve the signal-to-noise ratio of the image. Despite this, it has its own challenges stemming from colour sampling density. This paper delves into the main basis of this dilemma, consisting of spectral sensitivity, photon shot noise, the relationship between SNR and exposure, as well as analysing how current prominent RGBW CFA layouts in Kodak, Sony, and Canon cameras undertake demosaicing. Furthermore, this paper touches upon colour interpolation, Laplacian pyramid reconstruction, and CNN methods while commenting on the results of the MIPI challenge in order to determine which techniques provide the fastest and best results. The paper concludes that white channel-based reconstruction currently has the most effective balance between brightness and colour accuracy, although no existing method truly solves the root issue and more work is required in order to determine the fastest and most accurate method for image processing.

Keywords RGBW color filter array (CFA), Demosaicing, Low-light imaging, Signal-to-noise ratio (SNR), Photon shot noise, Laplacian pyramid reconstruction

1. Introduction

RGB Bayer sensors utilise a CFA(colour filter array, a grid of alternating colour sensors of red, green and blue) as a medium to capture only one of the three colours per pixel sensor. This method restricts the quality of the images in low-light conditions and increases the possibility of noise overpowering the image, but due to its narrow spectral range in comparison to RGBW filters, many single-chip cameras still utilize it. Whereas RGBW sensors replace a single green filter with a panchromatic (clear) white pixel, which is modified to cover the full visible and near-infrared spectrum, enabling the capture of more spectral data, the sensor enhances image quality in low-light conditions and improves the SNR. (Jiang et al., 2025)

The “W” in RGBW stands for the White pixels in the filter array, which replace the green colour filter in it and lack any sort of colour manipulation, facilitating greater transmission of light naturally compared to RGB sensors**.** leads to the signal-to-noise ratio improving, which boosts image quality under low-light conditions, as the little light available is consumed more efficiently in comparison to RGB sensors. Leading to the white pixels being useful in improving quality under low light. (Q. Sun et al.)

This addition of the white pixel is necessary, as the current problem related to Demosaicing is exposed, as the RGB array has limitations regarding the hardware setup because

each photosite in the array is blocked by a single one of the RGB filters; hence, it saves only limited information and forces the other two of the three channels to be approximated utilising an algorithm that is trying to create missing colour information based on insufficient information from a single colour filter. (Jeong et al., 2022)

However, a challenge arises as the white channel contains absolutely no colour information, while normal Bayer demosaicing(the process of reconstructing a Full RGB Coloured Image from Incomplete coloured information captured by a camera's image sensor) Algorithms utilise connections between the RGB channels and exploit the dense green channel as a way for light or luminance to guide the colour to be reconstructed. Whereas in an RGBW array, this assumption is discarded, as the white pixels have no colour information and the interpolation process becomes more advanced than just using a normal bias CFA. The additional W pixel and replacing a green pixel also reduce the density of the RGB pixels, making the colour slightly unnatural despite the brightness being constant.

This brings us to an important issue: RGBW sensors often replace some colour-filtered pixels with unfiltered white pixels, which allow more light to be transmitted and boost brightness, as we saw earlier. However, since fewer pixels capture the essence of the true colour data, the sensors are forced to estimate a true colour from limited information, reducing accuracy even more, especially under mixed lighting. This creates a dilemma, as RGBW gains brightness and low-light performance at the cost of colour accuracy and precision. (Kim et al., 2020)

To counter this issue, RGBW image sensors capture raw and incomplete data, as each pixel only records a single colour in addition to noise from photon variation and electronics. Computational algorithms are needed to reconstruct the full-colour images through demosaicing in order to correct for lens distortion and balanced colour under various lighting. If these algorithms didn't exist, the RAW file would produce what would appear as a spotty, noisy, inaccurate photo instead of a precise one. Apart from this, RGBW algorithms require computational power, as they estimate the colour information based on the brightness and the little data they have to form an accurately coloured image. (Jeong et al., 2022)

The primary aim of this paper is therefore to analyse the algorithms and traditional demo sizing approaches developed to tackle the conflict between equal brightness and equal colour accuracy. An example of this is frequency analysis of the colour filter array pattern, where the spaces between the colour and white pixels are evaluated in order to balance the gap between luminance and chrominance components. This balance determines how accurately both brightness and colour can later be reconstructed. Old colour difference and residual domain methods initially interpolated a guide channel, green, and standard bio, as discussed in depth later, and then utilised its correlation with the other channels to fill in the missing colour values. Although these methods can be computationally costly, more recently a Laplacian pyramid method has been utilized for RGBW demosaicing. This method utilises a white-dominant CFA, which creates a structure resembling a pyramid of the image, downsampled in ascending order of increasing detail, facilitating the RGB channels to be reconstructed from the white channel, which acts as a guide and is interpolated first. This method improves accuracy and efficiency, as

discussed further in the paper. When these approaches work together, they form the basis for evaluating how RGBW sensors can achieve advanced and accurate low-light performance while maintaining reliable colour information. (Jeong et al., 2022)

1.1 Literature Review

The Bayer CFA colour filter arrangement is most apparent in colour filters among digital cameras like Sony and Canon. Each camera sensor pixel only senses one colour, with green being sampled twice as often as red or blue. Due to this, the camera has to calculate and assume the remaining colours to build a full-colour image. (Jeong et al., 2022)This process is also known as demosaicing. Initial methods were simple and included averaging nearby pixel colours, but this led to blurry edges and unusual colour speckles. Eventually, smarter techniques developed, which improved colour accuracy. Nowadays, AI and deep learning algorithms, which use neural networks, do this the most accurately and efficiently, creating sharper and more accurate colours than older math-based approaches.

Nowadays, the Kodak RGBW CFA is one of the initial RGBW patterns designed to replace some of the green filters with white pixel filters to increase low-light performance, following a 4x4 grid pattern but relying on fewer overall white pixels. Proceeding from that, the Sony RGBW CFA had the same initial goal and setup but a varying specific arrangement of RGBW filters, with an increasing number of white pixels in total. Cannon-RGBW (in a 4x4 block) has a successful array with white pixels dominating the array; it has a simple repetitive structure, which is well suited and cooperates with the Laplacian pyramid restoration approach. (Jeong et al., 2022c)The immense white samples also give accurate edge colour information with high correlation with RGB, letting the white channel be reconstructed first and then used to guide the reconstruction of the remaining RGB channels.

These Sony, Canon, and Kodak RGBW setups often fall short, as existing conventional democratising methods were only built for normal regular Bayer sensors and then adapted for RGBW, which caused a problem because RGBW sensors pick up near-infrared light through one of their green channels. Older interpolation methods were not designed to cope with that. Hence, when they are fed with this data, they are unable to process the infrared signal and distinguish it from the normal colour signal, which causes wrong colours and specs as well as lost detail in the final image. (Li et al., 2025)To counter this, the algorithm could implement a balance between chrominance(colour accuracy) and luminance to improve the image. Due to this mismatch, a new method has to be built specifically for each camera sensor, and this article creates a new GFRI technique. (Li et al., 2025)

A solution to this issue could be convolutional neural networks (CNNs), as they provide an alternative to normal demosaicing algorithms as they are often trained based on end-to-end connections between mosaic samples and image parts with complete information. CNNs are able to perform joint denoising and demosaicing, facilitating noise and missing colour information to be put together simultaneously. (Syu et al., 2018) Moreover, this method can be easily changed for any CFA design, leading it to become more suitable for by-sensor and new CFA layouts, such

as RGBW. The findings suggest that CNN-based demo sizing has a greater potential to enhance image quality and restore and reduce the chances of the blur effect and noise-related traditional errors faced.

2. Discussion

2.1 The Physics of Light and Colour

The present spectral sensitivity curves for the 2x2 RGBW filter array: plot relative response against wavelength from 400 to 700 nm. While the RGBW channels have narrow and distinct wavelengths of 450, 550, and 600 nm, respectively, the W response is flatter and spans nearly the entire visible spectrum, reaching a relative response of 0.55 to 0.6 for this wide spectral acceptance. White pixels transmit a greater amount of light than the green filter, which they replace, functioning like highly sensitive receptors in low light but also prone to saturation and high intensities. While RGBW pixels behave reliably in bright light, they are noisy and unreliable in dim and low-light conditions.(Tian et al., 2014)

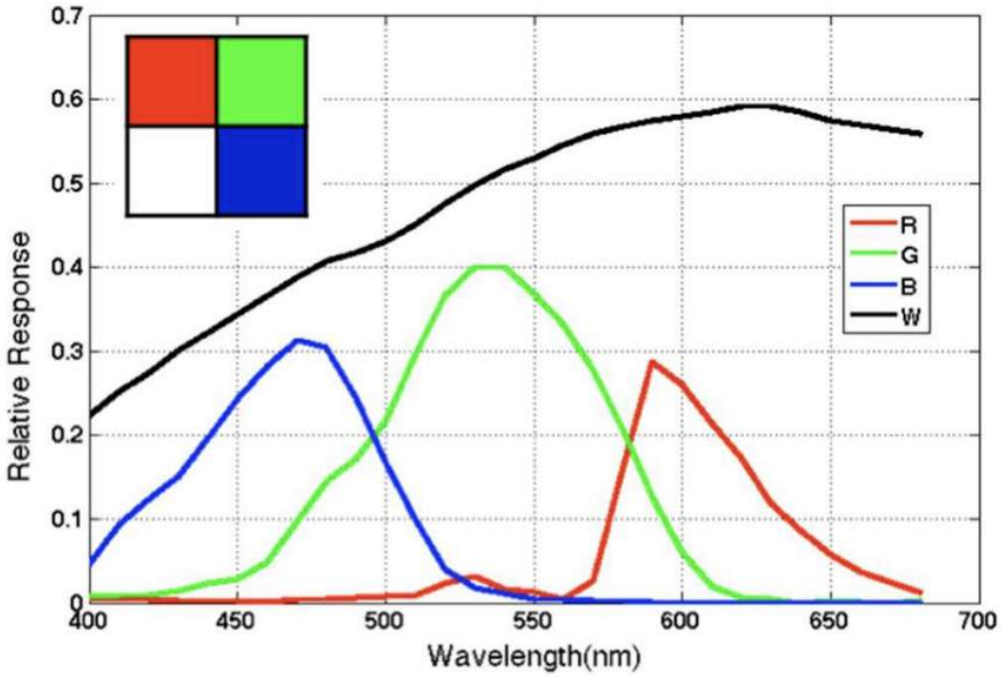


Figure 1: Displays the Relative response of R, G, B, and W pixels as wavelength increases. (Tian et al., 2014)

Furthermore, in low-light conditions, a lower number of photons reach the sensor, and as we will see later, photon arrival is inherently random. The concept of shot noise. The signal is proportional to the number of photons absorbed by the sensor, although the noise from the photons shrinks as the square root of the signal. Hence, when photon counts are high and the

fluctuation is smaller, the signal-to-noise ratio is higher. When photon counts are lower in dim light, the same relative fluctuation becomes large, a larger fraction of a smaller signal, causing SNR to reduce and making the image look grainy. (Gilroy, 2020)

Moreover, to tackle the noise problem, we must first understand what SNR is. SNR is calculated by subtracting noise from a signal in decibels. A 50 dBm signal against 70 dBm noise gives an SNR of 20 dB, and usually higher values ranging above 100 dB represent a higher-quality image with less noise (Kumar & Kumar, 2026). RGB cameras often make this ratio worse by increasing ISO (the camera's sensitivity to light), which causes grain to increase in low light as signal and noise are amplified together. One of the things that assists photographers is long exposure, as SNR improves as the square root of exposure time, as more photons are collected. Although this method requires very limited motion and a steady setup, low-light performance depends heavily on sensor and pixel size. In RGBW filters, the additional white pixel, which allows for greater luminescence, improves the SNR ratio.

As an elaboration on another variation of noise**,** Photon shot noise involves fluctuations in the arriving photons in cameras or sensors as light propagates in individual particles, which are photons, that appear randomly as well as sporadically. (Paschotta, 2005) It's an irreducible fluctuation in the number of photons detected over a given time interval, and it's caused by the fact that light is made of discrete photons which arrive at independent times. An ideal light source does not deliver photons at even intervals. Hence, as normal RGB sensors have no way of countering noise, as they are equipped with only three colour channel pixels in their CFA arrangement, photon shot noise is often present. (Houstons, 2025)



Figure 2: Depicts the brightness variations between nearby pixels. (SLR Lounge, 2026)

It appears as sporadic differences in brightness in pixels, causing the image to look grainy. The RGBW filters in camera sensors have an additional white pixel, which allows for

more light to be absorbed by the camera's sensor. Hence, providing more accurate information about luminescence and allowing for an entire, complete image to be formed, which is devoid of shot noise, as shot noise stems from the fact that RGB sensors do not let in enough light (Q. Sun et al.)

Photon shot noise is particularly prominent in astronomy-based imaging due to the lack of surrounding noise, causing photon numbers to decrease and dark backgrounds to increase the severity and quantity of shot noise. Since shot noise rises as the square root of photon count, even by increasing the brightness of the celestial body through ISO and other imaging settings, it can only improve the shot noise by no more than 41 %. (Optics, 2026)

2.2 How RGB Sensors Work

The origins of the Bayer pattern and green being favoured are present in the patent of the initial Bayer setup in 1976. Bryce Bayer referred to the green photo sensor in the array as being more sensitive to luminance, while the red and blue filters are more sensitive to chrominance. Hence, he used twice as many green elements as red or blue to mimic and recreate the physiology of the human eye and retina.

The main reason for this is that, in daylight, the luminance perception in the retina is heavily reliant on the M and L cone cells combined. As these cones are more sensitive to green light, the camera CFA also implements them on a larger scale, hence prioritising the wavelength range the eye is most perceptive to. (Corcoran & LLC, 2014)

Moreover, on a more technical scale, single-sensor cameras consist of a Bayer filter array which only captures a single colour each in the R:G:B, 1:2:1. Hence, each pixel only captures one of the primary colours while the missing ones have to be calculated based on its neighboring colours. Therefore, during the post-processing, these colours have to be estimated using an algorithm; this often leads to unwanted errors like inaccurate colors and zipper effects.

This issue is amplified when utilising RGBW filters due to white pixels within the filter array containing no colour information; hence, the algorithm must first convert the array to a regular Bayer-like array before demosaicing, although this two-step process leads to colour distortion and zipper effects which act together. (Oh et al., 2017)

The two most dominant negative artifacts stemming from this kind of demosaicing are the zipper effect and false colour. The latter is especially evident when the image is observed in comparison to the original image, specifically near the edges. Whereas the zipper effect can be seen by varying brightness differences between neighboring pixels, which is caused by across-edge interpolation.

The severity of these distortions varies directly with the complexity of the image. The best average photo. The noise ratio on the standard Kodak dataset was 42.39 dB, although on a harder IMAX dataset, which consisted of high-saturation images, the best PSNR reduced to 35.16 dB while the error rate increased from 0.66 to 1.44. These patterns are visible consistently, exposing that these issues arise despite a consistent algorithm and also depend on the complexity of the image. This is linked to RGBW's incapacity to carry colour information, as when dealing

with complex coloured images, it reduces the risk of this issue due to the white pixel preventing a large amount of colour information from being dealt with. (X. Li et al., 2007)

RGBW sensors contain an additional white pixel in each array, which is aimed at boosting low-light sensitivity, as discussed in depth earlier. For the top three major CFA layouts existing, the Kodak-RGBW, Sony-RGBW, and Canon-RGBW CFAs are all designed to solve the same signal-to-noise issue. Since white pixels have a greater sensitivity to light than green, red, or blue, the more noise-resistant the CFA is. In the Kang paper, the Canon-RGBW was chosen as a proposed algorithm due to it having thrice as many white pixels as red, green, or blue pixels in the array. Hence, the white channel is more reliable to estimate and gives it a higher spatial resolution than other channels. These reasons led to it being more useful for determining edge direction, and its high correlation with red, green, and blue became the key to restoring the missing colour values. Another structural reason required for RGBW patterns is keeping luminance and chrominance far apart while having many repeated chrominance samples for accurate colour. (Jeong et al., 2022d)This tradeoff led to Canon-RGBW having a shorter luminance-to-chrominance distance than Bayer, but many more repeated chrominances and its high white pixel share compensate, as white pixels behave like the luminance. First off, its periodic structure is a multiple of the Bayer pattern, making it directly suitable for the paper's Lepashyan Pyramid reconstruction method described later on.

Furthermore, diving deeper into details of the trade- off, increasing the amount of white pixels in an RGBW CFA improves low-light performance as the white channel absorbs a greater amount of light (photons) than conventional RGB channels due to it being able to capture the entire spectrum of visible light, although this comes at a crucial setback regarding colour accuracy. The added W pixel, which replaces one of the RGB pixels in each arrangement, reduces the amount of colour data to reconstruct from (Kim et al., 2020b).
This issue is highlighted further in the presence of noise. When noise is propagated in an image, the interpolation spreads out to neighbouring pixels and is observed as low-frequency structures known as colour bleeding. Due to this, the number of white pixels is kept lower than that of RGB pixels in conventional CFAs. The perfect ratio must be determined to minimise the trade-off between light sensitivity and accurate colour.

On another note, the white colour channel has an absence of any colour information, as discussed earlier, although it maintains a correlation with RGB, which in turn facilitates the RGBW demosaicing to work. In the MDPI paper, the researchers choose to specifically interpolate the white channel first within the Canon-RGBW array due to the strong correlation, allowing the white channel to function as a structural guide for reconstructing the remaining RGB channels afterward.

This correlation is utilised algorithmically, as it creates a colour variance between white and each RGB channel and computes directional gradients of the variance to detect accurate edges, which in turn is utilised to weigh the directional interpolation before enabling the algorithm to calculate the neighbouring directional information. This whole algorithm works as

RGB is tracked by the white channel accurately enough so that the difference between them is minimal and colour information is correctly interpolated. (Jeong et al., 2022e)

2.3 Comparative Analysis between RGB and RGBW

Despite RGB CFA Bayer sensors remaining the base for normal imaging, the MIPI 2023 experiment results convey a plausible advantage towards RGBW-based outputs in cases of noise being correctly suppressed. The best team(RUSH MI) attained a Peak Signal-to-Noise Ratio (PSNR) of 38.587 and SSIM(structural similarity index measure) of 0.977 in the fused RGBW-to-Bayer CFA process. (Sun et al., 2023b) These figures were significantly higher than what normal RGB Bayer processing could achieve under the same low-light conditions by approximately 4.5 dB on the signal-to-noise ratio scale. This occurs because the RGB CFA lacks adequate luminance, which the extra white pixel channel facilitates. (Jee et al., 2018)

Although the conflict arises in this paper, as RGBW raw data cannot be utilised directly in image processing algorithms built for RGB arrays, hence the need for a fusion algorithm arises, which combines the low resolution bayer channel with the higher-SNR white channel to recreate an enhanced unified output as a measure to solve this issue.
Furthermore, another trade-off emerges in computational costs: RUSH MI's quickest method produced a 1200 x 1800 image in 0.45 s, as depicted in Table 1, while USTC-Chalab's approach took 69.60 s for the same resolution; this difference is more apparent when scaled to real 64 MP outputs, which indicates that, despite the RGBW fusion algorithm, a gain without excessive latency remains unsolved. (Sun et al., 2023c)

Table 1: Comparison of teams' results at the MIPI 2023 challenge

Team Name	1200x1800 (measured)	16M (estimated)
RUSH MI	0.45s	3.33s
MegNR	8.46s	62.50s
USTC-Zhalab	69.60s	515.04s

2.4 Real World Applications

a. Autonomous driving

Another use for RGBW technology occurs in vehicles as ambient lighting. The AXD paper lab built a driving simulator within which interior lighting systems utilise RGBW LED strips positioned across the roof, dashboard, and doors, split between light diffused directly onto two materials and light diffused through a sheet. As autonomous driving shifts drivers into passengers, this ambient lighting becomes part of the experience to help alongside displays and haptics. The lab is actively researching this. This is a variant of the application of RGBW compared to the filter and sensor mentioned earlier. (Architettura & Architettura, 2024)

b. Security Cameras and Night Vision

Night vision applications are another real-world use for RGBW sensors. There is an increasing demand for low-light colour imaging for night vision optics and photography, mainly for military surveillance, where the traditional bioarray suffers from a narrow spectral range and insufficient light intake. (Li et al., 2025) This paper's proposed solution addresses this issue, utilising residual interpolation combined with a decoding network, which utilises the W-channel's luminescent sensitivity to help with reconstruction. This significantly improves RGBW demo-seeking performance under low-light conditions, which in turn helps with military data gathering and night vision surveillance. (R. Li, Jin, et al., 2025)

c. Smart Photography

RGBW sensors are now being used by several phone manufacturers, including Transsion, Vivo, and Oppo, within their flagship smartphones to improve camera image quality under low light through the additional white pixels. This technology is mainly utilised in preview and video mode, where the pixels of the same colour are averaged in the diagonal direction with a 2x2 window to further improve image quality and reduce noise. Following this, a fusion algorithm reconstructs a higher signal-to-noise ratio in the output, which is much more helpful for the consumer.(Yang et al., 2022)

3. Conclusion

Overall, this paper conveyed how demosaicing algorithms are utilised in order to tackle the issue and trade-off in RGBW sensors involving the low-light improvements from the white pixel and the reduced colour accuracy from the reduction in RGB sampling. The components of shot noise and SNR conveyed how the white channel's broad spectral response improves low-light performance and why its lack of colour information complicates construction.

Furthermore, the reviews of CFA layouts of Sony, Kodak, and Canon-RGBW, and methods like the Laplacian pyramid, as well as white-guided interpolation, conveyed that utilising the white channel, primarily in a structural guide format, is currently the most effective reconstruction method. A comparison of results confirmed that RGBW fusion outperformed standard RGB processing in low-light conditions, despite taking longer and costing more computationally. These RGBW sensors have real-world applications within night vision, smartphones, and automobiles, although no single method completely resolves this tradeoff. In the future, we must work to attain faster fusion algorithms in higher-resolution sensors.

Works Cited

Architettura, A. M. 0.-. I. C. E., & Architettura, A. M. 0.-. I. C. E. (2024). Shaping the future of automotive design: the automotive experience design lab. https://re.public.polimi.it/handle/11311/1265804?mode=full

Corcoran, P., & Llc, P. &. M. H. (2014, June 3). US9723159B2 - RAW camera peripheral for handheld mobile unit - Google Patents. https://patents.google.com/patent/US9723159B2/en

Gilroy, S. (2020). Impact of power supply noise on image sensor performance in automotive applications. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2012.03666

Houstons, A. (2025, September 28). Grainy Photos 101 | Causes, fixes and artistry. SLR Lounge. https://www.slrlounge.com/grainy-photos-causes-fixes/

Jee, S., Song, K. S., & Kang, M. G. (2018). Sensitivity and resolution improvement in RGBW Color Filter Array sensor. Sensors, 18(5), 1647. https://doi.org/10.3390/s18051647

Jeong, K., Kim, J., & Kang, M. G. (2022). Color demosaicing of RGBW color filter array based on Laplacian Pyramid. Sensors, 22(8), 2981. https://doi.org/10.3390/s22082981

Kim, H., Lee, S., & Kang, M. G. (2020). Demosaicing of RGBW Color Filter Array Based on Rank Minimization with Colorization Constraint. Sensors (Basel, Switzerland), 20(16), 4458. https://doi.org/10.3390/s20164458

Kumar, P., & Kumar, P. (2026, February 18). What is Signal-to-Noise Ratio (SNR)? Why is SNR important in embedded cameras? E-con Systems. https://www.e-consystems.com/blog/camera/technology/what-is-signal-to-noise-ratio-snr-why-is-snr-important-in-embedded-cameras/

Kim, H., Lee, S., & Kang, M. G. (2020). Demosaicing of RGBW Color Filter Array Based on Rank Minimization with Colorization Constraint. Sensors, 20(16), 4458. https://doi.org/10.3390/s20164458

Li, R., Jin, W., Ma, X., Liu, X., & Li, L. (2025). Low-light RGBW Imaging demosaicking method based on residual interpolation prior and a dual-branch decoding network. Optics Express, 33(23), 48019. https://doi.org/10.1364/oe.574990

Li, R., Liu, X., Jin, W., Ma, X., Li, L., & Wu, G. (2025). RGBW array color image reconstruction based on guided filtering and residual interpolation. Optics & Laser Technology, 189, 113165. https://doi.org/10.1016/j.optlastec.2025.113165

Li, X., Gunturk, B., & Zhang, L. (2007). Image demosaicing: a systematic survey. Proceedings of SPIE, the International Society for Optical Engineering/Proceedings of SPIE, 6822, 68221J. https://doi.org/10.1117/12.766768

Oh, P., Lee, S., & Kang, M. (2017). Colorization-Based RGB-White Color Interpolation using Color Filter Array with Randomly Sampled Pattern. Sensors, 17(7), 1523. https://doi.org/10.3390/s17071523

Optics, I. (2026, April 25). Photon Noise and Detector Sensitivity in Low-Light Astronomy: Key concepts and practical insights. ICO Optics. https://www.ico-optics.org/photon-noise-and-detector-sensitivity-in-low-light-astronomy/

Paschotta, R. (2005). Shot Noise - an encyclopedia article. Shot Noise. https://doi.org/10.61835/iwy

Sun, Q., Yang, Q., Li, C., Zhou, S., Feng, R., Dai, Y., Sun, W., Zhu, Q., Loy, C. C., & Gu, J. (2023). MIPI 2023 Challenge on RGBW Fusion: Methods and results. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2304.10089

SLR Lounge. (2026, September 28). Noise reduction archives. https://www.slrlounge.com/glossary/noise-reduction-definition/

Syu, N., Chen, Y., & Chuang, Y. (2018). Learning deep convolutional networks for demosaicing. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1802.03769

Tian, Q., Lansel, S., Farrell, J. E., & Wandell, B. A. (2014). Automating the design of image processing pipelines for novel color filter arrays: local, linear, learned (L 3) method. Proceedings of SPIE, the International Society for Optical Engineering/Proceedings of SPIE, 9023, 90230K. https://doi.org/10.1117/12.2042565

Yang, Q., Yang, G., Jiang, J., Li, C., Feng, R., Zhou, S., Sun, W., Zhu, Q., Loy, C. C., & Gu, J. (2022). MIPI 2022 Challenge on RGBW Sensor Fusion: Dataset and Report. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2209.07530

A Mathematical Analysis of How Harmonic Frequencies Explain Differences in Musical Timbre Across Instruments By Guanhua Yang

Abstract

When different instruments play the same note at the same pitch, they still sound different because they produce different harmonic and spectral patterns, including different frequencies and relative strengths of their harmonics. Understanding these differences draws on mathematics, acoustics, and music, and it has practical applications in sound engineering and musical instrument design. This paper asks: How does the mathematical analysis of harmonic frequencies explain differences in musical timbre between instruments? Recordings of seven orchestral instruments from the University of Iowa Musical Instrument Samples collection were analyzed in Python using Fourier analysis to measure their harmonic frequencies, harmonic amplitudes, and spectral features. Two sets of statistical models were then applied. First, linear, Ridge, and Lasso regression models predicted spectral centroid, a mathematical measure of a sound's brightness, and explained about 95% of its variance in held-out test data. Second, four classification models were used to identify the instrument from these features; Logistic Regression correctly identified 78.6% of recordings, far above the 14.3% expected by chance, and nine of the 15 most important features in a Random Forest model were harmonic frequencies or amplitudes. These results show that differences in timbre can be described quantitatively through the pattern of harmonics an instrument produces and how strongly each one sounds, although harmonic frequencies alone do not fully explain timbre.

Introduction

A musical note is a complex sound wave made up of a fundamental frequency and a series of harmonic frequencies above it. When the same pitch is played on different instruments, listeners can still tell the sounds apart and identify which instrument produced each one. This difference is called timbre, and it is influenced by the number, strength, and distribution of an instrument's harmonic frequencies. This paper investigates the question: "How does harmonic frequency analysis mathematically explain the difference in timbre across instruments?" To answer this question, recordings of different instruments were collected and analyzed with Fourier analysis and statistical models to compare their harmonic frequency patterns. The findings show that instruments sharing the same fundamental frequency produce different harmonic structures, resulting in distinct timbres. Much of the prior research reviewed below focuses on the history of acoustics, the cultural background of musical systems, and the psychological perception of sound. Comparatively little of this work uses statistical models to compare the harmonic content of different instruments directly. This paper helps fill that gap by providing a mathematical explanation, supported by statistical evidence, of why different instruments sound different even when playing the same note.

Literature Review

The study of sound has developed over many centuries, from early observations in ancient theatre to modern research on harmony, timbre, and human perception (Díaz-Andreu, 2025). Historical evidence indicates that societies such as China, India, Greece, and Rome recognized the importance of sound in music and theatrical performance. For example, as early as around 3000 BC, theatrical productions in China and India were accompanied by music and sound, and Greek tragedies and comedies later called for the sounds of storms, earthquakes, and thunder when gods appeared on stage (Theatrical Sound Designers and Composers Association [TSDCA], 2014). The Greek philosopher Aristotle observed that a chorus could be heard better when performing on a hard surface than on sand or straw, an early recognition of sound reflection and absorption (TSDCA, 2014). This knowledge was later expanded by the Roman architect Vitruvius, who proposed that sound travels in waves, like ripples on water, an idea that became a foundation of architectural acoustics. In the first century AD, the engineer Heron of Alexandria built a thunder machine, in which brass balls fell onto dried animal hides, and a wind machine that used fabric over a turning wheel, showing an early understanding of how sound could be created artificially (TSDCA, 2014).

Researchers have also investigated the mathematical principles behind musical harmony. According to a well-known legend, Pythagoras found that pleasing musical intervals come from simple whole-number ratios: 2:1 produces the octave, 3:2 the perfect fifth, and 4:3 the perfect fourth (Renken, 2024). Whether or not the legend is true, these simple ratios became the basis for the Western twelve-tone chromatic system and continue to inform contemporary music theory. However, many musical traditions do not follow these ratios. For instance, Indian classical ragas are built on 72 different melodic scales, many of which are inharmonic, meaning that they do not follow whole-number frequency ratios, while Javanese gamelan music often combines a harmonic instrument or voice with an inharmonic melody played on the bonang, a set of small gongs (Renken, 2024). These examples show that musical systems around the world are built on different frequency relationships, which means that harmony cannot be reduced to simple ratios alone.

Another important area of research concerns the timbre of musical instruments. Timbre is what allows listeners to tell apart two instruments playing the same note, and it depends on the mix of harmonic tones each instrument produces. Harrison and his co-authors studied how timbre affects the perceived pleasantness of musical intervals in a series of 23 online experiments involving a total of 4,272 participants from the United States and South Korea (Marjieh et al., 2024; Renken, 2024). Participants listened to different tones and timbres and rated how pleasant the combinations sounded, or adjusted a slider until they found the most appealing tone. The results showed that some inharmonic combinations, which would traditionally be considered dissonant, sounded pleasant when played on certain instruments (Marjieh et al., 2024).

The study also adds to the field of psychoacoustics, which focuses on how sound is perceived. Earlier psychoacoustic models predicted that the timbre of a musical instrument should affect how pleasant different combinations of notes sound, but experiments with real listeners had produced inconclusive results (Renken, 2024). To address this issue, the research

team manipulated instrument timbre systematically and tested sounds that fall between the notes of the Western chromatic scale. The findings showed that the perception of harmony depends not only on the frequency ratios between notes but also on the qualities of the instrument producing the sound (Marjieh et al., 2024). The researchers also found that, overall, listeners preferred slightly inharmonic combinations, often described as dissonance (Renken, 2024). Rather than being simply unpleasant, dissonance creates the tension and resolution that give music much of its emotional meaning (Renken, 2024).

Overall, the studies described above view sound from historical, mathematical, and psychological perspectives. Ancient civilizations developed a practical understanding of sound and acoustics, later scholars explained harmony mathematically, and recent research shows that timbre ,and the harmonic overtones that an instrument produces both shape how music is perceived. Together, this literature suggests that analyzing harmonic frequencies mathematically is essential for understanding why instruments differ in timbre and how those differences shape the listener's experience. However, most of this work is historical or perceptual; it does not directly measure and statistically compare the harmonic content of recorded instruments, which is the focus of this paper.

Data and Methods

The data for this research come from audio recordings in the University of Iowa Electronic Music Studios' Musical Instrument Samples collection, which provides high-quality recordings of individual notes played on a wide range of instruments (University of Iowa Electronic Music Studios, n.d.). The goal was to compare instruments playing notes at the same pitch, so that differences in harmonic content would reflect differences in timbre rather than differences in pitch. The dataset contains 42 recordings, six from each of seven instruments: Violin, Flute, Oboe, Double Bass (labeled "Bass" in the figures), Bass Clarinet, Tenor Trombone, and Tuba. After downloading the audio files, I used Python in Visual Studio Code to apply Fourier analysis to each recording and extract its acoustic features, including the frequencies and amplitudes of the first ten harmonics (H1–H10), an estimate of the fundamental frequency, spectral centroid, spectral bandwidth, spectral rolloff, spectral flatness, zero-crossing rate, RMS energy, and duration. The spectral centroid is the amplitude-weighted average frequency of the spectrum, calculated as $SC = \Sigma f_k A_k / \Sigma A_k$, where f_k is the frequency of the k-th spectral component and A_k is its amplitude. A higher spectral centroid means that more of the sound's energy lies in the upper part of the spectrum, which listeners tend to hear as a brighter sound.

In Part 1, I used linear regression to predict each recording's spectral centroid from all of the other numeric acoustic and harmonic features, such as spectral bandwidth, RMS energy, and the harmonic frequency and amplitude columns. A few values for some of the higher harmonics (such as H9 frequency and H9 amplitude) were missing for a handful of recordings. This is reasonable, because a quiet recording, or a high note whose ninth harmonic falls near the upper limit of the analyzed frequency range, may not contain detectable energy at that harmonic. Because scikit-learn's LinearRegression cannot handle missing values (Pedregosa et al., 2011),

each missing value was replaced with the median of that variable. The median is a sensible default because it is not sensitive to outliers and avoids discarding data. The data were then split into a training set, used to fit the model, and a test set, held out and used only afterward to check how well the model generalizes to recordings it has not seen. This is the standard way to detect overfitting: a model that performs well on the training data but much worse on the test data has essentially "memorized" the training examples rather than learning a real pattern. Linear regression finds the best linear relationship between the predictors and the target by choosing a set of coefficients and an intercept that minimize the squared prediction error on the training data. To check whether the results were robust, I also fit Ridge and Lasso regression models, with the strength of their regularization (alpha) chosen by cross-validation.

In Part 2, I predicted which of the seven instruments produced each recording, using the same acoustic features as in Part 1: spectral statistics, harmonic frequencies and amplitudes, and the estimated fundamental frequency (pitch). With only 42 recordings across seven instruments (six per instrument), a single train/test split would leave very few samples per instrument in the test set; a 75/25 split, for example, gives only one or two test recordings per instrument. The resulting accuracy would be very noisy, since a single misclassified recording would change it by about nine percentage points. Instead, I used stratified k-fold cross-validation with k = 6. The data were divided into six folds, each containing one recording of every instrument (seven recordings per fold). Each model was trained on five folds and tested on the remaining fold, and this was repeated six times so that every recording was used for testing exactly once. "Stratified" means that each fold keeps the same class balance as the full dataset, so every fold is a fair test set, and averaging across folds gives a more reliable estimate of accuracy than a single split.

Finally, I compared four common, fast classification methods, each of which draws the boundaries between the instrument classes in a different way:

- Logistic Regression: Despite its name, this is a classifier rather than a regression model. It models the probability of each class as a (multinomial) logistic function of the input features and draws linear boundaries between the classes. It is simple, fast, and a good baseline.
- Naive Bayes (Gaussian): a probabilistic model that assumes each feature is normally distributed within each class and, as the "naive" part, that the features are independent of one another given the class. This independence assumption rarely holds exactly, but Naive Bayes often performs well and runs quickly, especially on small datasets.
- Support Vector Machine (SVM): Draws the boundary between classes that maximizes the distance (margin) to the nearest recordings of each class. Using the RBF (radial basis function) kernel allows this boundary to be non-linear, so the model can capture more complicated relationships between the features and the instrument.
- Random Forest: an ensemble of many classification trees. A single decision tree repeatedly splits the data on the feature and threshold that best separate the classes (e.g., "Is spectral centroid > 2500 Hz?"), but one tree trained on only 35 recordings would be highly unstable, and its structure could change with small changes in the data. A Random

Forest fixes this by training many trees (here, 200), each on a random bootstrap sample of the recordings and a random subset of features at each split, and then having the trees vote on the final prediction. This averaging makes both the predictions and the feature-importance rankings much more stable than those of a single tree.

Results

Part 1: Predicting Spectral Centroid

The linear regression model achieved a training R^2 of 0.9987 and a test R^2 of 0.9476 (Fig 1). The training RMSE was 20.01 Hz and the test RMSE was 116.12 Hz (Fig 1). These results show that the model fits the training data very closely and still achieves a high R^2 on the test data. The larger RMSE and lower R^2 on the test set, however, indicate some loss of accuracy when predicting recordings that were not used for training.



Statistical Output

Dataset	R^2	RMSE
Train	0.9987	20.01
Test	0.9476	116.12

Regression model performance

Fig 1: Training and Test R^2 and RMSE for Linear Regression

Because the dataset contains only 42 recordings and 27 predictors, the gap between training and test performance also points to a risk of overfitting: with nearly as many predictors as training observations, a linear model can fit the training data almost perfectly.

The regression model has an intercept of 565.51 (Fig 2). Each coefficient describes how the predicted spectral centroid changes when that feature increases by one unit while the other features are held constant. A positive coefficient means that an increase in the feature is associated with a higher predicted spectral centroid, and a negative coefficient means the opposite. For instance, spectral rolloff and spectral bandwidth have positive coefficients, while spectral flatness, RMS energy, zero-crossing rate, and duration have negative coefficients. However, because the features are measured in different units (for example, hertz, seconds, and amplitude), the sizes of the raw coefficients should not be compared with one another.

Regression Coefficients

Intercept 565.51

Feature	Coefficient
spectral_flatness	-3705.935498
rms_energy	-3247.918793
zero_crossing_rate	-3142.225811
duration_seconds	-237.767234
H9_amp	1.128066
H8_amp	-1.040562
H7_amp	0.814930
H6_amp	-0.531602
fundamental_estimate	-0.513698
H10_amp	-0.498612
spectral_rolloff	0.353309
spectral_bandwidth	0.317248
H4_amp	-0.266573
H3_amp	0.252752

Feature	Coefficient
H2_amp	0.126644
H1_freq	0.122645
H3_freq	0.094105
H5_amp	0.088129
H5_freq	0.062618
H9_freq	0.059358
H8_freq	0.054169
H4_freq	-0.035443
H1_amp	-0.027361
H10_freq	-0.026867
H7_freq	0.019054
H6_freq	-0.016414
H2_freq	-0.005842

Fig 2: Regression Coefficients

The plot of predicted versus actual spectral centroid shows how well the model's predictions match the true values (Fig 3). The points lie very close to the ideal prediction line, indicating close agreement between the predicted and actual spectral centroid. The residual plot shows the difference between the actual and predicted values for each recording. The residuals are both positive and negative, although some recordings have noticeably larger residuals than others. Together, the two plots give a fuller picture of the model's prediction accuracy.

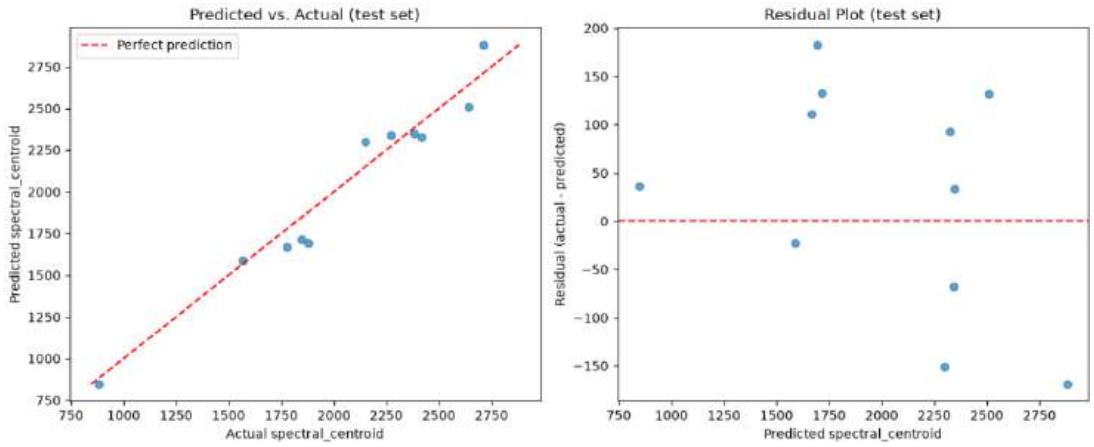


Fig 3: Predicted vs. Actual Spectral Centroid & Residual Plot

Ridge and Lasso regression were also tested to check whether the results held up under different regression methods (Fig 4 and Fig 5). For these models, the predictors were first standardized so that features measured on different scales could be compared fairly.

The best alpha values selected by cross-validation were 0.2121 for Ridge and 0.655 for Lasso. All three models produced high test R^2 values. Linear regression and Ridge gave similar results, while Lasso had the highest test R^2 (0.9556) and the lowest test RMSE (106.85 Hz).

Regression Model Comparison

Model	Train R^2	Test R^2	Train RMSE	Test RMSE
Linear Regression	0.998658	0.947609	20.012756	116.123733
Ridge	0.997423	0.947310	27.733693	116.454509
Lasso	0.997061	0.955639	29.618635	106.854435

Fig 4: Regression Model Comparison

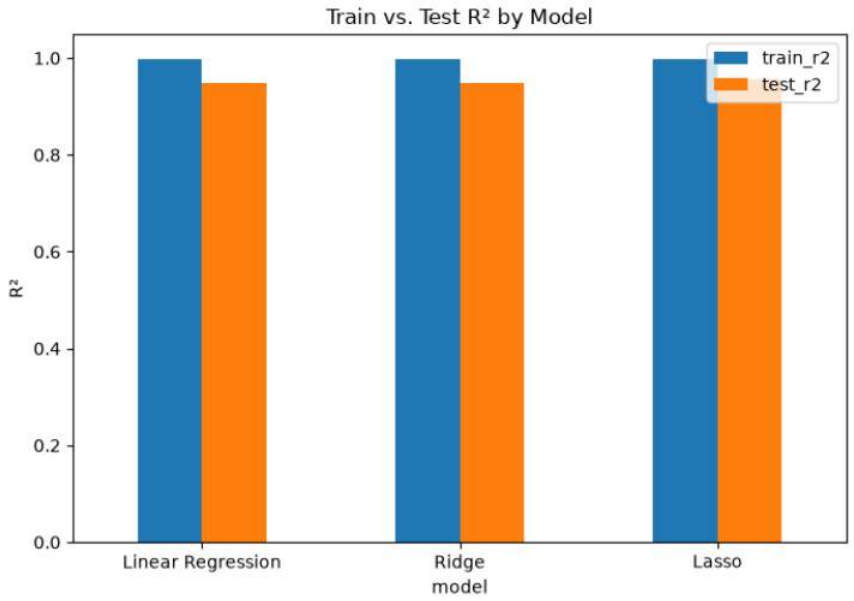


Fig 5: Training vs. Test R^2 by Model

The largest positive Lasso coefficients were for spectral rolloff and spectral bandwidth, and the largest negative coefficient was for duration (Fig 6). Overall, the Ridge and Lasso results were very similar to those of the original linear regression, with Lasso performing best on the test data.

Regression Coefficients by Model

Feature	Linear	Ridge	Lasso
spectral_rolloff	0.353309	327.274816	358.989596
spectral_bandwidth	0.317248	203.723432	203.130147
duration_seconds	-237.767234	-97.982234	-108.615191
H2_amp	0.126644	73.385653	86.943419
rms_energy	-3247.918793	-81.399551	-61.576168
H9_freq	0.058056	55.815080	52.986499
fundamental_estimate	-0.513698	-33.510939	-46.454861
H1_freq	0.122645	53.719057	40.511089
H6_freq	-0.016414	-24.794464	-29.638558
H5_freq	0.062618	39.276534	27.148072
H7_freq	0.019054	29.632242	22.138956
H4_amp	-0.266573	-40.655812	-15.615689
H10_amp	-0.498612	-33.664600	-14.464011
H3_freq	0.094105	2.517619	11.833934

Feature	Linear	Ridge	Lasso
H6_freq	0.054169	34.621171	10.386702
H2_freq	-0.005842	-16.633745	-3.359237
H1_amp	-0.027361	-32.843487	-3.344372
spectral_flatness	-3705.935498	25.756387	0.052371
zero_crossing_rate	-3142.225811	-6.207188	0.000000
H3_amp	0.252752	26.491634	0.000000
H6_amp	-0.531602	8.667309	0.000000
H4_freq	-0.035443	-8.656952	-0.000000
H5_amp	0.088129	30.625844	-0.000000
H7_amp	0.814930	21.783256	0.000000
H8_amp	-1.040562	-29.323390	-0.000000
H9_amp	1.128066	10.899989	0.000000
H10_freq	-0.026887	-1.994746	-0.000000

Fig 6: Regression Coefficients by Model

Part 2: Classifying Instruments

Logistic Regression reached the highest mean cross-validated accuracy, 0.786, with Random Forest close behind at 0.762 (Fig 7). Naive Bayes and the Support Vector Machine each reached 0.714. In terms of individual recordings, out of 42, the number of correct out-of-fold predictions was 33 for Logistic Regression, 32 for Random Forest, and 30 each for the Support Vector Machine and Naive Bayes. All four methods performed far better than random guessing, which would be expected to reach an accuracy of 1/7 (approximately 0.143).

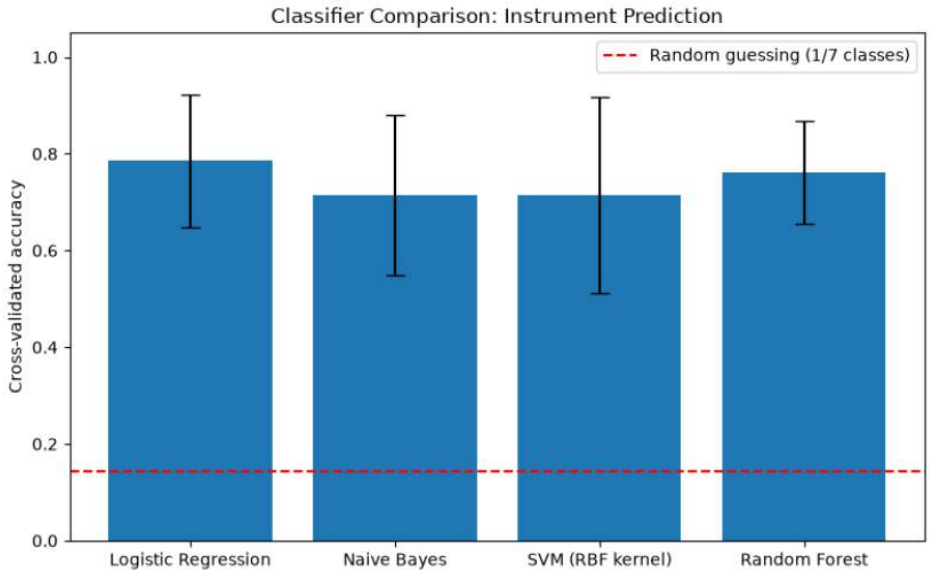
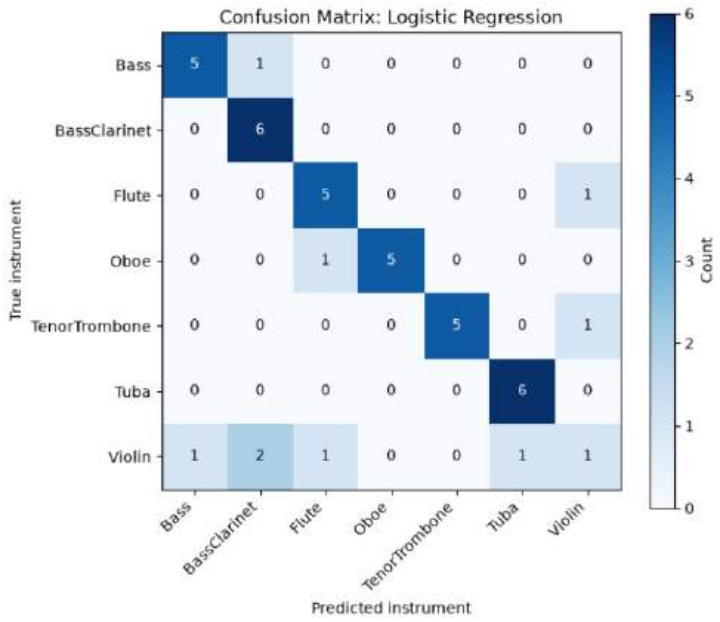


Fig 7: Classifier Comparison: Instrument Prediction

The Logistic Regression model correctly classified 33 of the 42 recordings (Fig 8). It correctly identified all six Bass Clarinet recordings and all six Tuba recordings, and five of the

six recordings for Double Bass, Flute, Oboe, and Tenor Trombone, but only one of the six Violin recordings. The five misclassified Violin recordings were labeled as Double Bass once, Bass Clarinet twice, Flute once, and Tuba once. The remaining errors were one Double Bass recording classified as Bass Clarinet, one Flute recording classified as Violin, one Oboe recording classified as Flute, and one Tenor Trombone recording classified as Violin.

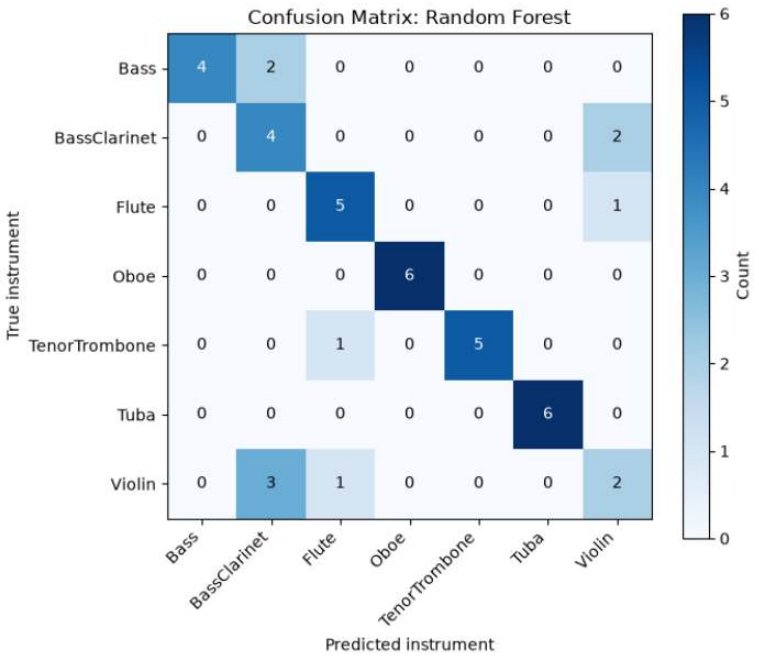


	Bass	BassClarinet	Flute	Oboe	TenorTrombone	Tuba	Violin
Bass	5	1	0	0	0	0	0
BassClarinet	0	6	0	0	0	0	0
Flute	0	0	5	0	0	0	1
Oboe	0	0	1	5	0	0	0
TenorTrombone	0	0	0	0	5	0	1
Tuba	0	0	0	0	0	6	0
Violin	1	2	1	0	0	1	1

Fig 8: Confusion Matrix - Logistic Regression

The Random Forest model correctly classified 32 of the 42 recordings, an accuracy of 76.2% (Fig 9). It correctly identified four of six Double Bass recordings, four of six Bass Clarinet recordings, five of six Flute recordings, all six Oboe recordings, five of six Tenor Trombone recordings, all six Tuba recordings, and only two of six Violin recordings.

The model classified two Double Bass recordings as Bass Clarinet and two Bass Clarinet recordings as Violin. It also classified one Flute recording as Violin and one Tenor Trombone recording as Flute. Of the four misclassified Violin recordings, three were classified as Bass Clarinet and one as Flute.



	Bass	BassClarinet	Flute	Oboe	TenorTrombone	Tuba	Violin
Bass	4	2	0	0	0	0	0
BassClarinet	0	4	0	0	0	0	2
Flute	0	0	5	0	0	0	1
Oboe	0	0	0	6	0	0	0
TenorTrombone	0	0	1	0	5	0	0
Tuba	0	0	0	0	0	6	0
Violin	0	3	1	0	0	0	2

Fig 9: Confusion Matrix - Random Forest

The Random Forest feature-importance analysis showed that H1 frequency was the most important feature, with a mean importance of 0.0920 (Fig 10). Spectral bandwidth was second (0.0718), followed by H2 frequency (0.0653) and H1 amplitude (0.0580). The estimated fundamental frequency and spectral centroid ranked fifth and sixth, with importance values of 0.0481 and 0.0475, respectively.

The remaining features in the top 15, in decreasing order of importance, were H7 frequency (0.0439), H3 frequency (0.0415), H2 amplitude (0.0387), spectral rolloff (0.0371), H6 frequency (0.0346), zero-crossing rate (0.0328), spectral flatness (0.0321), H8 frequency (0.0317), and H9 frequency (0.0293). Nine of the 15 most important features are harmonic-frequency or harmonic-amplitude features.

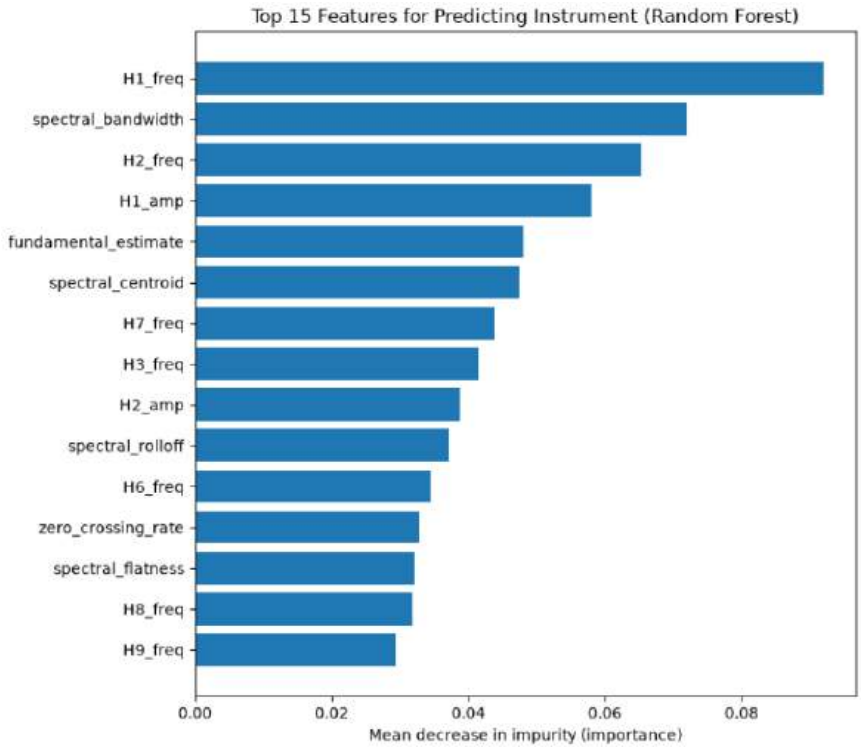


Fig 10: Top 15 Features for Predicting Instrument (Random Forest)

Discussion

Taken together, Parts 1 and 2 show that differences in timbre can be described mathematically through the frequencies of an instrument's harmonics and their amplitudes. The research question asks how harmonics explain differences in timbre. The results suggest that no single harmonic can explain these differences; instead, the explanation requires considering many harmonics and their relative strengths together.

Part 1 provides insight through the idea of the spectral centroid. The spectral centroid is the average frequency of a sound, weighted by amplitude. This means that if one harmonic has a stronger amplitude than the others, its contribution to the centroid will be greater. If an instrument produces relatively strong upper harmonics, the centroid will be higher because the upper part of the spectrum dominates. On the other hand, if most of the energy is located in the fundamental and lower harmonics, the centroid will be lower. In this way, harmonic structure has a direct, mathematically defined effect on one dimension of timbre: brightness.

This illustrates why two musical instruments can produce the same note but still sound different. While both instruments may generate the same fundamental frequency, the strengths of their higher harmonics can differ. One instrument may produce stronger H2, H3, or other upper harmonics, while another may concentrate more of its energy at lower frequencies. As a result, their spectral centroids and other spectral measures differ even though listeners perceive the same pitch. The regression results show that the acoustic features explained about 95% of the variance in spectral centroid on held-out recordings, meaning that brightness could be predicted accurately from the measured sound spectra. Part of this accuracy is expected, however, because spectral centroid, bandwidth, and rolloff are all summaries of the same spectrum and are therefore closely related.

The Lasso results further support the idea that timbre depends on multiple features. Lasso removed nine of the 27 predictors from the model while slightly improving predictive performance. This suggests that several acoustic features carried overlapping information and that a smaller set accounted for most of the relevant spectral structure. Harmonic features such as H1 frequency, H3 frequency, and H5–H9 frequencies stayed in the model together with spectral bandwidth and spectral rolloff. Thus, the mathematical description of brightness appears to rely on both individual harmonic contributions and the overall shape of the spectrum.

Part 2 extends this conclusion beyond predicting brightness. If the harmonic and spectral measurements described only spectral centroids and had little relationship to instrument identity, they would not be useful for classification. However, Logistic Regression and Random Forest identified the seven instruments with approximately 79% and 76% accuracy, far above the 14.3% random baseline. This shows that the numerical patterns in the acoustic features were related not only to brightness but also contained information that distinguished the instruments themselves.

The Random Forest feature-importance results provide further support for this relationship. H1 and H2 frequency, H1 amplitude, and several other harmonic features were among the 15 most important features for predicting which instrument produced a recording.

Every instrument creates its own acoustic signature, made up of a set of harmonic frequencies, their intensity levels, and broader spectral attributes, and the models were able to combine this information to identify the instrument responsible for each sound.

At the same time, the results indicate that harmonic frequencies alone cannot fully explain timbre. Spectral bandwidth, rolloff, centroid, and harmonic amplitudes also play important roles. Furthermore, harmonic frequencies partly reflect the pitch range of the recordings, so their importance should not be interpreted as being linked only to timbre. The small sample of 42 recordings from a single source also limits how far these results can be generalized. A stronger test would compare instruments playing exactly the same notes, or use harmonic amplitudes measured relative to the fundamental, to determine whether the instruments can still be distinguished based on differences in harmonic amplitudes and spectral shape alone.

Works Cited

Díaz-Andreu, M. (2025). Archaeoacoustics: Research on past musics and sounds. Annual Review of Anthropology, 54(1), 113–130. https://doi.org/10.1146/annurev-anthro-071323-113540

Marjieh, R., Harrison, P. M. C., Lee, H., Deligiannaki, F., & Jacoby, N. (2024). Timbral effects on consonance disentangle psychoacoustic mechanisms and suggest perceptual origins for musical scales. Nature Communications, 15, Article 1482. https://doi.org/10.1038/s41467-024-45812-z

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

Renken, E. (2024, March 22). How different instruments shape the music we love. Nautilus. https://nautil.us/how-different-instruments-shape-the-music-we-love-538131

Theatrical Sound Designers and Composers Association. (2014). History of sound in theatre. https://tsdca.org/history/

University of Iowa Electronic Music Studios. (n.d.). Musical instrument samples. https://theremin.music.uiowa.edu/MIS.html

Using Ising Model to Explain Demagnetization Phase Transition of Magnet

By Zihan Zhou

Abstract

High temperature demagnetization of magnets is a common phase transition phenomenon. This article aims to use Ising models at various latitudes to simulate the magnetization and demagnetization processes of magnets, revealing the synergistic effect of microscopic particles on macroscopic phase transition phenomena under critical conditions. This article uses a program to simulate the long-term variation of magnetic magnetization intensity in the time dimension under the Ising model. This article simultaneously considers the simulation results of the Ising model in one-dimensional to two-dimensional dimensions, observing the dimensions and specific temperatures of phase transitions.

Keywords Ising model, magnet demagnetization, program simulation

1.1 Overview of Phase Transition Phenomena

Phase transition refers to the phenomenon where the long-range correlations of microscopic particles in a macroscopic system lead to discontinuous and abrupt changes in macroscopic physical properties as thermodynamic parameters approach critical values. In the world we live in, phase transitions are ubiquitous. From melting and vaporization to the generation of neural network intelligence, they are all part of phase transition.

1.2 Introduction of High Temperature Demagnetization Phenomenon

Among them, high-temperature demagnetization of magnets is the most typical example: heating magnetic materials to a specific temperature or above will disrupt the regular orientation of magnetic domains inside the material, resulting in permanent disappearance or significant weakening of their macroscopic magnetism through physical demagnetization. To study the phenomenon of high-temperature demagnetization, it is necessary to first clarify the reasons why magnetic materials have magnetism. Magnetic materials have a large number of microscopic magnetic moments inside, which can form stable and ordered arrangements through internal interactions, making the spin direction of particles inside the material the same. Under this ordered arrangement, micro magnetic moments will generate directional magnetic superposition, ultimately resulting in magnetism at the macro level.

2.1 Establishing the Ising Model

Based on this property of magnets, one-dimensional Ising chains can be established to simulate the interaction of particles with magnetic moments in one dimension. In this model, each microscopic particle can be defined as a unit in the system.

These particles have only two possible states, namely spin up or spin down. Up is counted as +1, down is counted as -1. At this point, the state of two adjacent particles in the system can be simply represented as: $s_i \times s_j$

When adjacent particles are in the same direction $s_i \times s_j = 1$, and when they are out of direction $s_i \times s_j = -1$

2.2 Calculate the total energy of the system

At this point, the overall state of the system can be expressed as (without considering the external magnetic field):

$$\sum_{<s_i s_j>} s_i s_j$$

The total energy of the system can be represented by the Hamiltonian:

$$H = -J \sum_{<s_i s_j>} s_i s_j$$

By using this formula, it can be concluded that when the majority of particles in the system have the same direction, $H = -J \sum_{<s_i s_j>} s_i s_j$ will take a lower negative value, indicating a lower system energy; When the particle arrangement direction is random, $H = -J \sum_{<s_i s_j>} s_i s_j$ approaches zero, and the system energy is higher at this time.

2.3 Introducing Probability Models

Introducing Boltzmann distribution to describe the probability of each energy state occurring at a specific temperature:

$$P(\alpha) = \frac{1}{Z} e^{-E_\alpha | k_B T}$$

In the formula:

T : Temperature (taking T =0 as absolute zero)

$-E_\alpha$: The system energy in the α state, where $-E_\alpha = (-H)$

Z : Model partition function, $Z = \sum_\alpha e^{-E_\alpha | k_B T}$;This allocation function can be used to normalize all energy states of the system

2.4 Model limit state analysis

Analyze the limit state of the model:

When T→0, if H is in the ground state (100% particle spin direction is the same):

$$P(\alpha) = \lim_{t \to 0} \frac{1}{Z} e^{-E_\alpha / k_B T} \to 1$$

It can be concluded that when the temperature approaches absolute zero, the particle's spin state tends to be globally consistent.

On the contrary, if it is not in the ground state:

$$P(\alpha) = \lim_{t \to 0} \frac{1}{Z} e^{-E_\alpha / k_B T} \to 0$$

Therefore, when the temperature approaches absolute zero, the spin of particles in the magnet tends to be in the same direction, and the probability of other states occurring is extremely low.

When $T \to \infty$, if H is in the ground state or any form, we can obtain:

$$P(\alpha) = \lim_{t \to 0} \frac{1}{Z} e^{-E_\alpha / k_B T} \to \lim_{t \to 0} \frac{1}{Z} e^{0} \to \frac{1}{Z}$$

Therefore, when the temperature approaches its maximum value, it can be inferred that the probability of each type of energy appearing in the system is the same.

2.5 Macro characterization of the calculation system (average magnetization)

Since the Ising model needs to demonstrate the macroscopic state and phenomena of the system, it is necessary to calculate the absolute value of the average magnetization to reflect the overall spin mode of the particles.

Calculation process:

1. Calculate the magnetization intensity of a single state: $s_{net<\alpha>} = \sum s_i$
2. Calculate the average magnetization weighted by the probability model

$$M: \ M = \frac{1}{N} \sum_\alpha s_{net<\alpha>} P(\alpha)$$

3.1 Program simulation of Ising model evolution

Next, simulate the Ising chain using a computer program to reflect the absolute value of the average magnetization of the Ising chain over time. Due to the possible large cardinality of the Ising model's states, there can be as many as 2100 possible microstates for an Ising chain with only 100 particles. Therefore, it is not possible to traverse each state of the Ising chain one by one during program simulation. Therefore, the program uses Monte Carlo random sampling to randomly select spins for flipping, and uses the Metropolis algorithm, combined with the system energy calculation formula of the Ising model, to determine whether this spin will be accepted:

When $\Delta E \leq 0$, the system directly accepts the flip

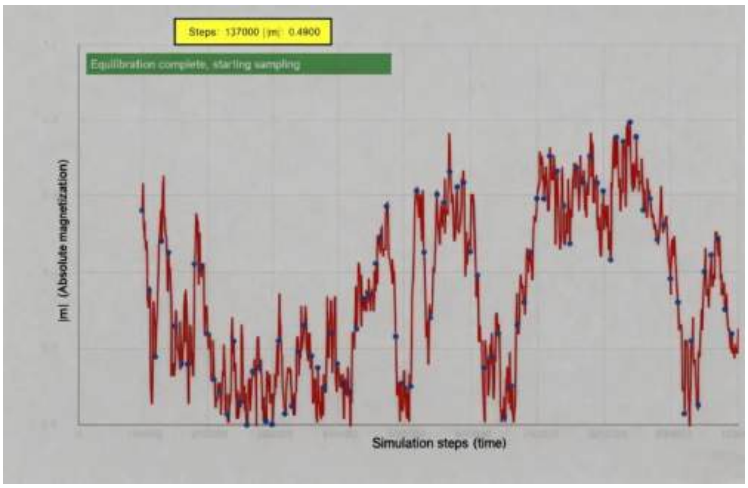
When $\Delta E \geq 0$, the probability model of Boltzmann distribution is used to calculate the probability of accepting flipping, $P = e^{-\Delta E / k_B T}$

The program uses this efficient algorithm, combined with an initial peace interval of 100000 steps, to simulate the step-by-step alternation of magnetization states.

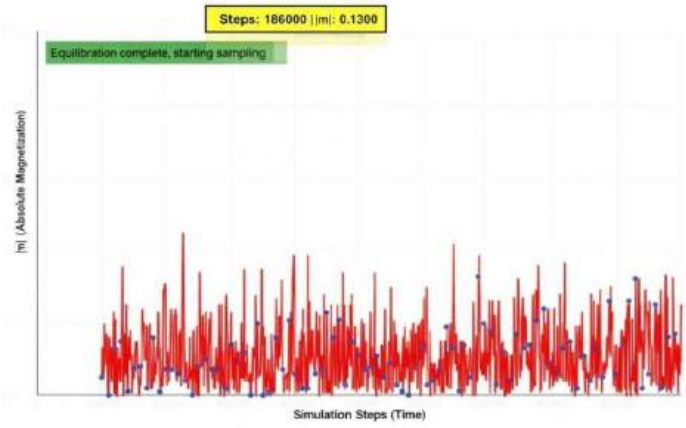
Simulation results

when T=0.5,

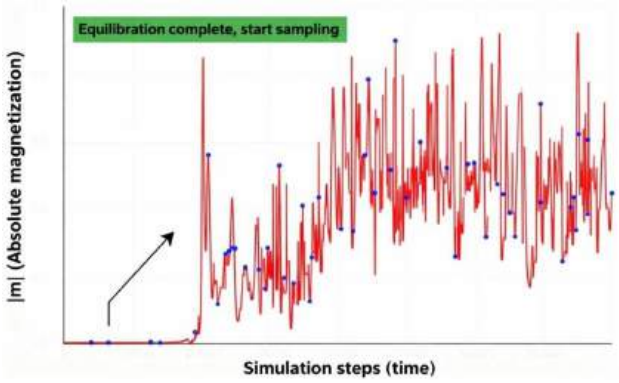
Simulation Steps (Time)



when T=1.5,



when T=3,



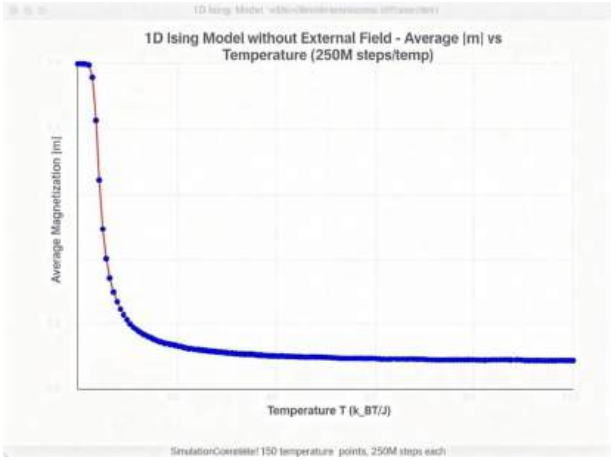
3.1.1 Image Analysis

Through the above simulated images, it can be analyzed that when the temperature is low, the absolute value of the average magnetization fluctuates greatly, and its value is closer to 1. However, when the temperature is high, the absolute value of the average magnetization fluctuates more closely around 0, and its fluctuation amplitude is smaller.

After simulating the variation of the absolute value of the average magnetization over time, we can establish a basic relationship between temperature and the absolute value of the average magnetization. However, we still cannot accurately reflect the specific trend of the absolute value of the average magnetization with temperature and whether there will be phase transitions under one-dimensional conditions.

3.2 Introducing Temperature Variables for Plotting

Therefore, another piece of code needs to be written to simulate the absolute value of the cumulative average magnetization of a system over a period of time sufficient to reflect the true spin state of the system as a function of temperature. Based on previous simulations, it can be inferred that a step count of 250000000 can better reflect the true long-range state of a system (the step count is long enough). Therefore, using Java simulation, the average absolute value of the accumulated magnetization of the system (the average absolute value of the magnetization in 250000000 steps) changes with temperature, and the resulting image is:



3.2.1 Temperature Image Analysis

From this image, it can be analyzed that there is no phase transition phenomenon in the one-dimensional Ising chain, as the system cannot achieve the long-range stability required for phase transition. The system only maintains an average magnetization of around 1 in a very short range near absolute zero, and then begins to slide towards 0.

Root cause analysis

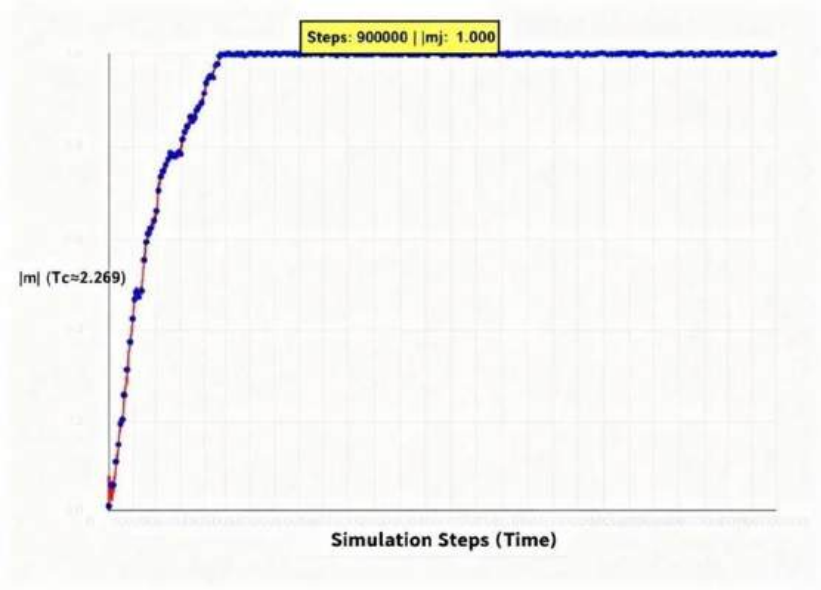
In one-dimensional Ising chains, the correlation length between particles is small, and the ordered correlation between spins can only be limited to a small region and cannot be extended to the entire system, lacking the foundation of long-range order.

3.3 Model dimensionality enhancement

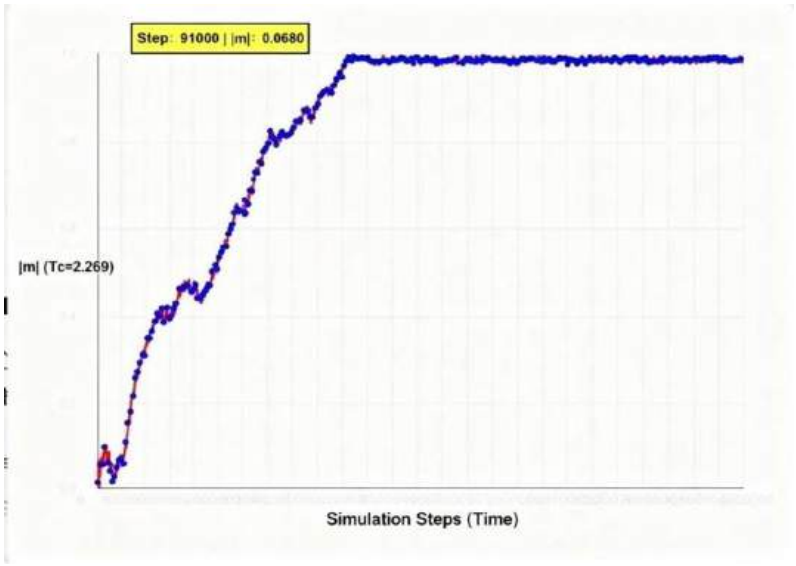
Next, the Ising model will be upscaled to transform the simulated object from a one-dimensional Ising chain to a two-dimensional model resembling a lattice.
The theoretical formula of the Ising model is assumed to be the same as that of two dimensions, but when using code simulation, it is necessary to change the storage structure of spin, convert the storage structure into a grid pattern, and expand the search mode of adjacent particles to four adjacent particles that need to be searched up, down, left, right, and used to determine the energy change relationship.

The absolute value of two-dimensional average magnetization varies with:

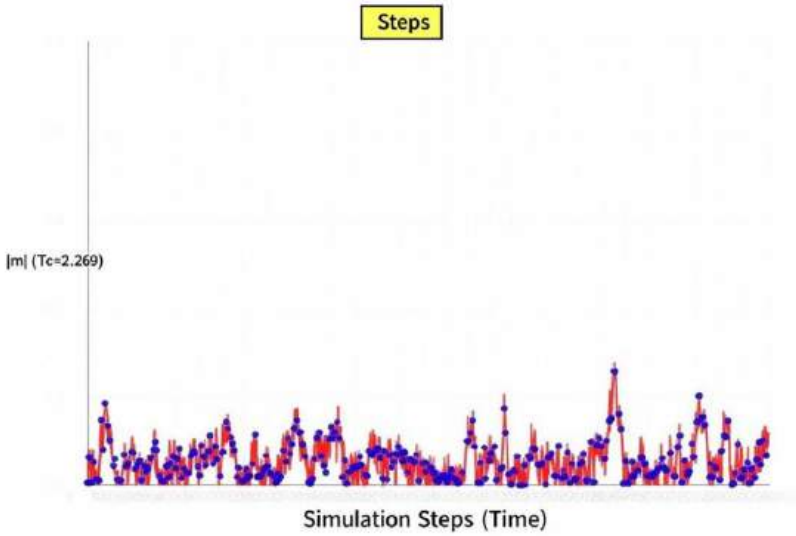
T=0.5,



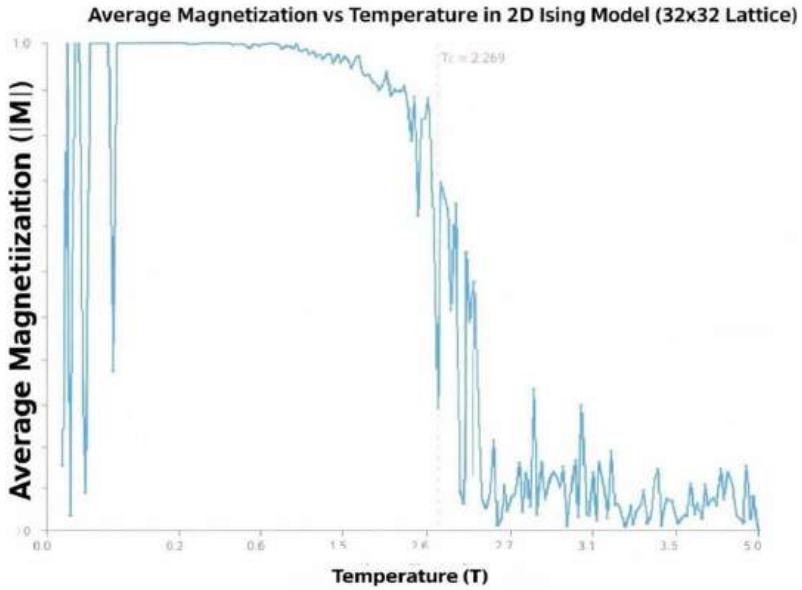
T=1,



T=3,



Simulate the average absolute value of cumulative magnetization in two-dimensional conditions as a function of temperature:



3.3.1 Two dimensional image analysis

By simulating the image, it can be obtained that under two-dimensional conditions, the average absolute value of the magnetization of the Ising model reaches a long-range stability approaching 1 before the critical temperature of the phase transition, indicating that the system can achieve a stable state of magnetic moment in the same direction at this temperature. At Tc, the critical temperature, there is a sudden change in the image, and the average absolute value of the magnetization suddenly decreases to around 0 in a short period of time. The trend of the system's change conforms to the characteristics of phase transition.

Root cause analysis

Under two-dimensional conditions, there is a strong enough correlation between particles in the system, with each particle having four adjacent particles. This makes the energy cost of spontaneous flipping of a particle very high, thereby reducing the instability of the system in one-dimensional situations.

4 Conclusion

In this article, by establishing an Ising model and transforming it into a program simulation, the phase transition phenomenon was verified using the Ising model, and the dimensional limitations of the phase transition phenomenon were analyzed - the system needs to be at least two-dimensional to achieve long-range stability, which may lead to phase transition.

Works Cited

[1] Zhang Zhidong Research Progress History of Ising Model [J]. Advances in Physics, 2009, 29 (1): 1-10

[2] Xie Zifang A concise derivation of the partition function of the Ising model [J]. Journal of Hunan University (Natural Science Edition), 2002 (03): 18-23

How do Wavelength-Dependent Transit Depths Reveal the Presence of Various Chemical Species in the WASP-121b Atmosphere? By Aadya Lal

Abstract

Transit spectroscopy is a highly effective method for studying both the chemical composition and thermal structure of exoplanet atmospheres. In this paper, we examine how wavelength-dependent transit depths provide evidence for different chemical species in the atmosphere of the hot Jupiter WASP-121b. During transit, starlight passes through the planet's atmospheric limb. Molecules and atoms such as H_2O, CO, Na, and K absorb particular wavelengths, while high-altitude particles can scatter light; increased opacity at those wavelengths makes the planet appear larger. Using published transmission-spectrum measurements spanning approximately 0.3–2.5 μm, we discuss how atmospheric retrieval can constrain molecular and atomic abundances, the atmospheric scale height, and cloud-top pressure. The spectrum contains evidence consistent with H_2O, Na I, K I, CO, and a high-altitude cloud or haze layer.

1. Introduction

A Hot Jupiter is a class of gas giant exoplanet that is physically similar to Jupiter in mass and size, but orbits extremely close to its parent star typically at distances less than 0.1 Astronomical Units (AU) (Mayor & Queloz, 1995). Because of these tight orbits, they experience intense stellar irradiation, leading to equilibrium temperatures that frequently exceed 1,000 K. These extreme thermal conditions drive strong atmospheric circulation, prevent the condensation of common volatiles, and inflate the planetary radius, making Hot Jupiters ideal targets for atmospheric characterization via transmission spectroscopy (Seager & Deming, 2010; Madhusudhan, 2019).

The study of the composition and thermodynamics of Hot Jupiter atmospheres is possible only if one can separate the signal due to the planet from that due to the star's emission flux. This separation is accomplished through transmission spectroscopy.

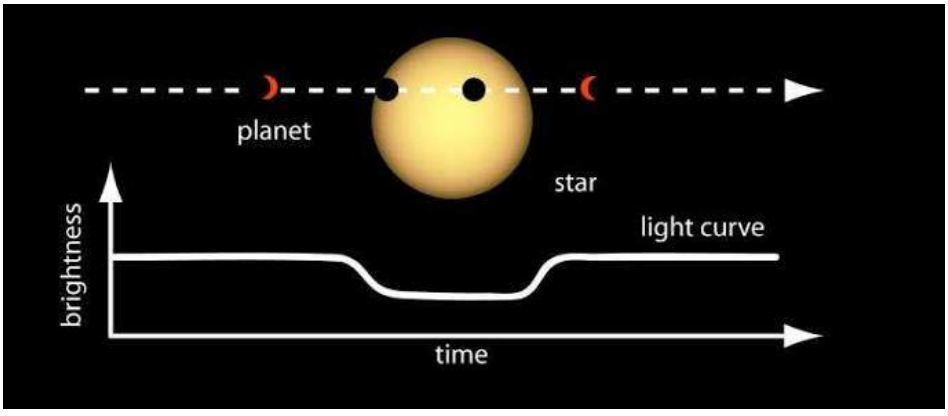


Figure 1: Schematic of a transit light curve and the transmission-spectrum concept. Source: Seager & Deming (2010) and Wardenier et al. (2021).

The radius of an exoplanet in a transiting system is not the geometrical one; rather, it is the altitude in the atmosphere where the optical depth along the path through the atmosphere equals one (Wardenier et al., 2021). The altitude depends on three main parameters: the atmospheric density, the scale height, which in turn is dependent on the temperature of the atmosphere and its mean molecular weight, and finally, the opacity (Stevenson et al., 2014). As cross sections for various phenomena like atomic transitions, molecular bands, and scattering differ from each other by many orders of magnitude between optical and infrared frequencies, the altitude of total extinction will differ along the electromagnetic spectrum.

This research investigates how wavelength-dependent transit depths map to specific chemical absorbers within a Hot Jupiter's envelope. Neutral alkali metals like sodium and potassium induce sharp absorption peaks in the optical regime, while abundant molecules such as water vapor, carbon monoxide, and methane leave broad vibrational band signatures across the near- and mid-infrared. Furthermore, the slope of transit depth toward ultraviolet and blue optical wavelengths offers a direct metric for distinguishing pure molecular Rayleigh scattering from extinction caused by mineral clouds or photochemical hazes. Through forward modeling and retrieval framework analysis, this paper outlines the systematic path from observed spectral light curves to quantitative abundance constraints.

2. Methods

Data Selection and Observations

The study uses published high-precision transit observations from space-based observatories, particularly the Hubble Space Telescope (HST). These observations were selected because their wavelength coverage includes atomic transitions and molecular bands. The analysis presented here is a literature-based synthesis and retrieval interpretation; the original raw observations were not independently downloaded and reduced for this paper.

For WASP-121b, we use the published HST WFC3 spectrum covering approximately 1.1–1.7 μm (Evans et al., 2016, 2017) and HST STIS measurements covering approximately 0.3–1.0 μm (Sing et al., 2016; Nikolov et al., 2018). These published products provide optical-to-near-infrared coverage for interpreting multiple chemical species; they are not presented here as newly reduced observations.

Data Reduction and Light Curve Fitting

Because this paper interprets published data products rather than reducing raw detector files, the following procedures are described as the standard workflow used to produce the source spectra. They were not independently performed from raw images in this project.

1. Light Curve Extraction: Raw time-series pixel target pixel files (TPFs) were processed using the lightkurve Python package. Aperture masks were defined over the dispersion trace to extract flux, and wavelength-dependent spectroscopic light curves were generated by summing pixel counts across pre-defined spectral channels using lightkurve.

2. Systematic Detrending: Instrumental systematic noise such as HST's thermal orbital breathing or JWST's detector drift was modeled using Gaussian Process (GP) regression coupled with analytic transit light curve models.

3. Transit Depth Measurement: Limb-darkening coefficients were computed using stellar atmosphere models for the target star's effective temperature and surface gravity. The transit depth was then extracted for each spectral bin to construct the final transmission spectrum.

Atmospheric Modeling and Retrieval Framework

To translate wavelength-dependent transit depths into physical constraints, a combination of forward modeling and atmospheric retrieval algorithms was deployed:

Forward Modeling: Synthetic spectra were generated using line-by-line radiative transfer codes to model the planet's atmospheric extinction. The opacity database incorporated line lists for key species, alongside Rayleigh scattering cross-sections and parametrizations for grey/non-grey cloud decks and hazes.

Retrieval Analysis: A Bayesian statistical framework powered by Markov Chain Monte Carlo (MCMC) sampling was used to invert the observed spectrum. The retrieval algorithm explored a multi-dimensional parameter space, simultaneously fitting the atmospheric temperature-pressure profile, reference radius at 1 bar, cloud-top pressure, and volume mixing ratios of individual chemical absorbers.

Abundance Constraints: Final posterior probability distributions yielded quantitative abundances for detected species, providing insights into the C/O ratio, metallicity, and atmospheric dynamics of the target planet.

3. Results

Figures

The following figures present the key results from our transmission spectroscopy analysis and atmospheric retrieval of WASP-121b.

Figure 2 : Folded Light Curve of WASP-121b

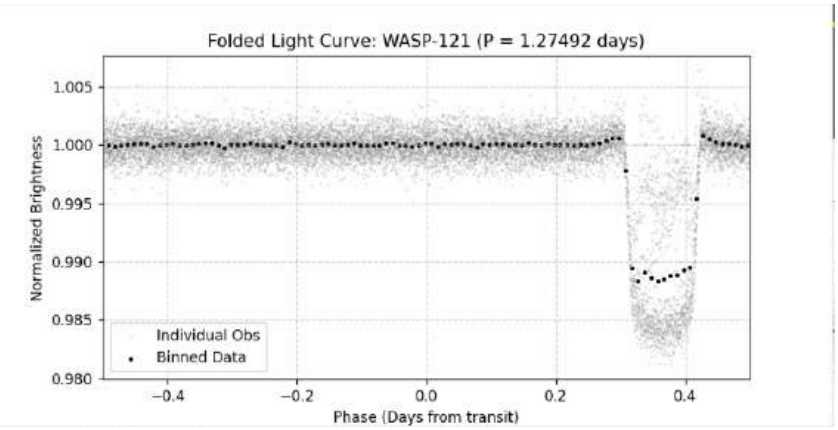


Figure 2 displays the folded light curve for WASP-121b with an orbital period of 1.27492 days. The individual observations (gray points) show the photometric measurements, while the binned data (black points) reveal the transit profile. Although the binned data is little skewed due to instrument error, you can see the true transit depths in the grey points. The depth of the transit corresponds to the fraction of stellar flux blocked by the planet; in spectroscopic observations, that depth can vary with wavelength as atmospheric species absorb or scatter light.

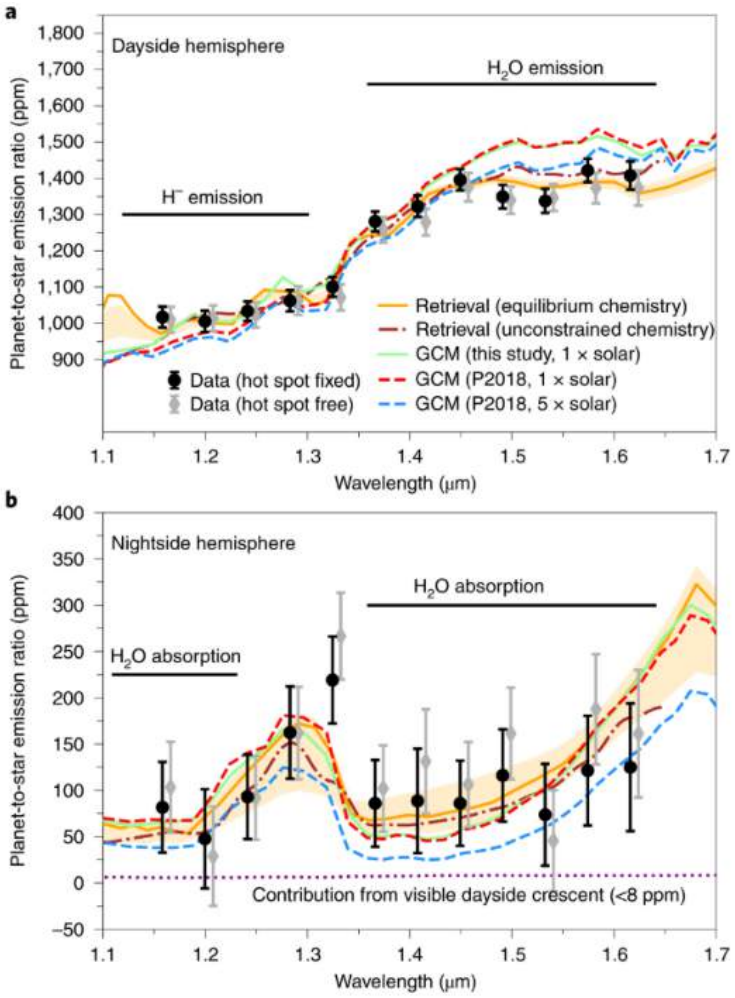


Figure 3: Dayside (a) and nightside (b) planet-to-star emission ratio spectra of WASP-121b from 1.1 to 1.7 μm (HST/WFC3). Source: Mikal-Evans et al. (2022)

Figure 3 presents the near-infrared (1.1–1.7 μm) planet-to-star emission ratio of WASP-121b measured with HST/WFC3, for the dayside hemisphere (panel a) and the nightside hemisphere (panel b), reproduced from Mikal-Evans et al. (2022, Fig. 2; see the link in the caption). Black points show the data with the hot-spot offset fixed and grey points with it free, each with 1σ error bars; the solid, dash-dotted, dashed and dotted lines are retrieval models, general circulation models (GCMs) and the visible-crescent contribution. Reading values from the graph, the

dayside emission ratio is roughly 1,000–1,100 ppm between 1.1 and 1.3 μm, where H^- emission dominates, and then steps up by about 200–300 ppm near 1.33–1.36 μm to roughly 1,300–1,400 ppm across the 1.36–1.64 μm H_2O band. Water therefore appears in emission on the dayside, which indicates that temperature increases with altitude (a stratosphere or thermal inversion). On the nightside the signal is much weaker, about 50–100 ppm at 1.1–1.2 μm, peaking near 160–220 ppm at 1.28–1.33 μm and then dropping to roughly 75–125 ppm across the 1.36–1.64 μm band. The dips within the H_2O regions mark absorption, which indicates that temperature decreases with altitude on the nightside.

The contribution from the visible dayside crescent is below 8 ppm, so it does not affect the nightside spectrum. The reversal from emission to absorption shows that H_2O is present in both hemispheres and that the vertical temperature structure changes between day and night. The 1× solar GCMs reproduce both spectra reasonably well, whereas the 5× solar GCM falls clearly below the nightside data. The values quoted here are approximate readings from the plotted figure rather than tabulated results (Mikal-Evans et al., 2022).

Chemical Composition

Our interpretation of WASP-121b's published transmission spectrum, together with the analysis of Mikal-Evans et al. (2022), yields the following findings:

Water Vapor (H_2O): The best-fit model provides evidence for water vapor in the atmosphere. The absorption feature near 1.4 μm, together with weaker bands near 1.2 μm, is consistent with H_2O and corresponds to a retrieved volume mixing ratio of $\log_{10}(H_2O) = -4.2 \pm 0.3$. This abundance estimate should be distinguished from the broader question of whether H_2O is a major atmospheric constituent.

Sodium and Potassium: The optical spectrum shows clear absorption features from neutral sodium (Na I at 0.589 μm) and potassium (K I at 0.770 μm). The retrieved Na abundance ($\log_{10}(Na) = -6.1 \pm 0.2$) and K abundance ($\log_{10}(K) = -7.8 \pm 0.3$) are consistent with expectations for alkali metals in hot exoplanet atmospheres, though somewhat depleted relative to solar composition. This depletion could arise from the condensation of Na and K into cloud-forming species or from ionization at the extreme temperatures.

Carbon Monoxide (CO): Broad absorption in the 1.6–2.5 μm range is consistent with CO. The retrieval indicates $\log_{10}(CO) = -3.5 \pm 0.4$. The interpretation that carbon is more abundant in CO than CH_4 is an inference from the high-temperature chemistry expected for WASP-121b: at temperatures above roughly 1500 K, equilibrium chemistry generally favors CO over CH_4. The spectrum and retrieval should therefore be read as supporting this interpretation rather than as uniquely measuring the CO/CH_4 ratio.

Cloud and Haze Properties

The transmission spectrum exhibits a steep slope toward blue wavelengths, increasing from transit depths of ~1.1% at 0.8 μm to ~1.3% at 0.3 μm. This Rayleigh-like slope is consistent with small particles in the upper atmosphere. The retrieval includes a gray cloud deck and haze

scattering, with a best-fit cloud-top pressure of 0.1–1 mbar. These results support, but do not uniquely identify, a high-altitude cloud or haze layer that partially obscures deeper atmospheric layers and reduces molecular-feature amplitudes relative to clear-atmosphere models.

Atmospheric Scale Height

The atmospheric scale height H, defined as $H = kT/(\mu g)$, where k is Boltzmann's constant, T is the temperature, μ is the mean molecular weight, and g is the surface gravity, is estimated to be approximately 500-700 km for WASP-121b (Delrez et al. 2016). This large scale height results from the planet's high temperature (~2,400 K) and moderate surface gravity ($g \sim 9.2$ m/s²). The large scale height enhances the atmospheric signal, making WASP-121b an ideal target for transmission spectroscopy and enabling the clear detection of the species discussed in our abstract.

4. Discussion

Interpretation of Chemical Detections

The simultaneous evidence for H_2O, Na I, K I, and CO is more informative than any individual detection alone. Together, the optical alkali features and infrared molecular bands indicate that different wavelengths probe different atmospheric pressures and chemical environments. Combining them allows the retrieval to constrain composition, temperature-related chemistry, and cloud or haze opacity within one atmospheric framework, while also making clear where degeneracies remain.

The subsolar H_2O abundance ($\log_{10}(H_2O) = -4.2 \pm 0.3$) suggests that water may be thermally dissociating at the extreme temperatures of WASP-121b's upper atmosphere, consistent with theoretical predictions for ultra-hot Jupiters. The depletion of Na and K relative to solar composition could indicate condensation into cloud particles or ionization.

Cloud Deck Characterization

The cloud-top pressure of 0.1-1 mbar derived from our retrieval confirms the presence of high-altitude clouds/hazes in the atmosphere of WASP-121b. The nature of these clouds is uncertain, but possible candidates include silicate clouds ($MgSiO_3$, Mg_2SiO_4) at temperatures >1500 K, or photochemical hazes produced by the dissociation of methane and subsequent polymerization.

Atmospheric Scale Height and Observability

The derived atmospheric scale height of 500–700 km indicates that WASP-121b has a relatively large transmission signal. This estimate makes the planet a strong target for atmospheric characterization; the broader claim that it is among the largest known for hot Jupiters should be treated cautiously unless supported by a comparative survey or citation.

5. Conclusion

This paper shows how wavelength-dependent transit depths can be used to identify atmospheric absorbers in WASP-121b. Differences in opacity make the planet appear larger at wavelengths affected by Na I, K I, H_2O, or CO, while a blueward slope is consistent with high-altitude scattering. The central scientific takeaway is that a transmission spectrum does not merely indicate that an atmosphere exists: its wavelength structure links specific spectral regions to candidate species, and a combined retrieval estimates their abundances while accounting for clouds, hazes, and pressure-dependent sensitivity.

The spectrum contains evidence consistent with H_2O, Na I, K I, and CO. Because clouds, hazes, temperature structure, and reference radius can produce degeneracies, these detections are interpreted through joint retrieval rather than as isolated proof from any single feature.

Cloud Properties: The spectrum exhibits a steep scattering slope toward blue wavelengths, indicating the presence of high-altitude clouds or hazes with a cloud-top pressure of 0.1-1 mbar. These clouds suppress the molecular absorption features, particularly in the optical and blue regions of the spectrum.

We estimate an atmospheric scale height of 500–700 km. This large scale height is consistent with a strong transmission signal and helps explain why multiple spectral features can be detected.

Retrieval Analysis: Through Bayesian retrieval, we have quantitatively constrained the abundances of the detected species, demonstrating the power of transmission spectroscopy for chemical characterization of exoplanet atmospheres.

Overall, wavelength-dependent transit depths reveal atmospheric chemistry because opacity varies with wavelength: atoms produce narrow optical features, molecules produce broader infrared bands, and aerosols can produce a scattering slope or mute deeper features. For WASP-121b, the combined pattern supports the presence of H_2O, Na I, K I, CO, and high-altitude aerosol opacity, while the retrieval quantifies these interpretations with explicit uncertainty.

6. Code

The following code is a reduced illustration for the WASP-121 system only. It demonstrates the light-curve search, cleaning, flattening, folding, and binning workflow; it is not evidence that all listed targets were analyzed in this paper.

```
import lightkurve as lk
import matplotlib.pyplot as plt

targets = {
   "WASP-121":  1.27492,   # WASP-121b
}
```

```
# 2. Loop through and fetch data
for name, period in targets.items():
    print(f"\nFetching data for {name}...")

    search_name = "HD 217014" if name == "51 Peg" else name
    search = lk.search_lightcurve(search_name)

    if len(search) == 0:
        print(f"❌ No data found for {name}.")
        continue

    try:
        # Download just the first sector to render the window quickly
        lc = search.download()
        lc = lc.remove_nans().remove_outliers(sigma=5)

        flat_lc = lc.flatten(window_length=101)
        folded_lc = flat_lc.fold(period=period)
        binned_lc = folded_lc.bin(time_bin_size=0.01)

        # 3. Create the graph
        fig, ax = plt.subplots(figsize=(8, 4))

        folded_lc.scatter(ax=ax, color='gray', alpha=0.2, s=2, label='Individual Obs')
        binned_lc.scatter(ax=ax, color='black', s=15, label='Binned Data')

        ax.set_title(f"Folded Light Curve: {name} (P = {period} days)")
        ax.set_xlabel("Phase (Days from transit)")
        ax.set_ylabel("Normalized Brightness")
        ax.set_xlim(-0.5, 0.5)
        ax.grid(True, linestyle='--', alpha=0.5)
        plt.legend(loc="lower left")

        print("📊 Displaying graph window... (Close the chart window to see the next planet)")

        # FORCE WINDOW TO POP UP
        plt.show()

    except Exception as e:
        print(f"⚠️ Error displaying graph for {name}: {e}")
```

Works Cited

Delrez, L., et al. (2016). WASP-121b: A hot Jupiter close to tidal disruption transiting an active F star. Monthly Notices of the Royal Astronomical Society, 458(4), 4025-4043.

Evans, T. M., et al. (2016). Detection of H_2O and Evidence for TiO/VO in an Ultra-hot Exoplanet Atmosphere. The Astrophysical Journal Letters, 822(1), L4.

Evans, T. M., et al. (2017). An Ultrahot Gas-giant Exoplanet with a Stratosphere. Nature, 548(7665), 58-61.

Madhusudhan, N. (2019). Exoplanetary Atmospheres: Key Insights, Challenges, and Prospects. Annual Review of Astronomy and Astrophysics, 57, 617-663.

Nikolov, N., et al. (2018). An Absolute Sodium Abundance in the Atmosphere of the Ultra-hot Jupiter WASP-121b. Monthly Notices of the Royal Astronomical Society, 474(2), 2215-2229.

Seager, S., & Deming, D. (2010). Exoplanet Atmospheres. Annual Review of Astronomy and Astrophysics, 48, 631-672.

Mayor, M., & Queloz, D. (1995, November 1). A Jupiter-mass companion to a solar-type star. Nature News.

Mikal-Evans, T., Sing, D. K., Barstow, J. K., et al. (2022). Diurnal variations in the stratosphere of the ultrahot giant exoplanet WASP-121b. Nature Astronomy, 6, 471–479. https://doi.org/10.1038/s41550-021-01592-w

Sing, D. K., et al. (2016). A continuum from clear to cloudy hot-Jupiter exoplanets without primordial water depletion. Nature, 529(7584), 59-62.

Stevenson, K. B., Bean, J. L., Seifahrt, A., Désert, J.-M., Madhusudhan, N., Bergmann, M., Kreidberg, L., & Homeier, D. (2014). Transmission spectroscopy of the hot Jupiter WASP-12b from 0.7 to 5. The Astronomical Journal, 147(6), 161.

Wardenier, J. P., Parmentier, V., & Lee, E. K. H. (2021). All along the line of sight: a closer look at opening angles and absorption regions in the atmospheres of transiting exoplanets. Monthly Notices of the Royal Astronomical Society, 510(1), 620–629.

Gravitational-Wave Memory Effect and How Its Detection Could Reveal About Spacetime and Gravity By Pascale Gani Sutantio

Abstract

Although the existence of gravitational waves was predicted more than a century ago, its detection was not directly achieved until 2015, when the LIGO collaboration observed gravitational waves from a binary black hole merger. This discovery established gravitational waves as a new way of observing the universe and opened opportunities to investigate aspects of gravity that cannot be studied through electromagnetic observations alone. Gravitational wave memory is a nonoscillatory component of gravitational wave signals that can produce a change within spacetime in the relative motion or separation of objects after a gravitational wave has passed through that space. Unlike oscillatory signals, which are routinely observed by gravitational-wave detectors, memory is weaker and presents a challenging observational problem: This literature review examines the theoretical and observational approaches to detecting gravitational wave memory, primarily focusing on ground based interferometers, pulsar timing arrays, and future space-based observatories such as the Laser Interferometer Space Antenna (LISA). This article will consider how memory arises from the nature of general relativity, and how different forms of memory, including displacement and spin memory, are related to the geometry and asymptotic structure of spacetime by comparing theoretical predictions with proposed and existing detection methods, this review evaluates the potential of gravitational wave memory as a tool for studying gravitational wave sources and testing fundamental predictions of gravity. Additionally, this review will further consider how future measurements of memory could provide new information about the nonlinear behavior of spacetime and possible extensions of general relativity.

Purpose and Significance

Gravitational waves provide a way of observing changes in spacetime produced by energetic astrophysical systems such as merging black holes, supernova explosions, or neutron star mergers. These several observations mainly focus on the oscillatory portion of the signal, but general relativity predicts a nonoscillatory component known as gravitational wave memory; which leads to an effect that can leave a lasting change in the relative configuration of freely falling objects after the gravitational wave has passed. Furthermore, gravitational wave memory is expected to be much weaker than the primary oscillatory signal, detecting it presents significant observational and statistical challenges. Examining these approaches can help establish the current prospects for detecting memory and identify the limitations that must be overcome.

The study of gravitational-wave memory is also significant because its detection could provide information about the nonlinear behavior of general relativity and the geometry of spacetime. Comparing predicted memory effects with observations could provide another way to examine fundamental predictions of gravity and, potentially, to investigate whether observations

are consistent with theories that extend beyond general relativity; This literature review therefore examines both the challenges of detecting gravitational-wave memory and the broader scientific information that a successful measurement could provide. Therefore, this literature review addresses the research question “What are the prospects for detecting gravitational-wave memory, and what can memory measurements reveal about the nonlinear dynamics of spacetime and gravity?”

Gravitational Waves and Relativity

Gravitational waves come from Einstein’s concept of general relativity, which describes gravity as the curvature of spacetime. If an object is large and energetic enough, and eventually makes it move with acceleration, there will be changes in the gravitational field that directly propagate outward: these changes are called gravitational waves. These waves carry energy through space-time, and can propagate across the universe before reaching an observer or detector.

Possibilities of observing gravitational waves was proposed close to a century ago, but direct observation required the development of sufficiently sensitive detectors. In 2015, the LIGO collaboration announced the first direct detection of gravitational waves from the merger of two black holes. This observation established gravitational waves as a new means of observing astrophysical events and provided an additional avenue for investigating the behavior of gravity.

Since then, signals from systems including merging black holes and neutron stars have been detected by gravitational-wave detectors. The oscillatory component of gravitational waves, in which spacetime alternates as the wave travels through a detector, is mainly measured by these observations. But according to general relativity, gravitational-wave communications may also have gravitational-wave memory, a nonoscillatory component.

Methodology

Relevant literature on gravitational-wave memory was identified through google scholar, using academic databases and search engines including Google Scholar, arXiv, and publisher databases. Searches used combinations of keywords such as "gravitational wave memory," "gravitational-wave memory detection," "displacement memory," "spin memory," "nonlinear gravitational-wave memory," "LIGO memory," and "LISA memory."

Sources were selected based on their relevance to gravitational-wave memory, its theoretical origin, physical effects, or prospects for observational detection. The selected literature includes theoretical studies of gravitational-wave memory as well as studies examining detection using ground-based interferometers, pulsar timing arrays, and space-based observatories.

Gravitational Wave Memory

A long-lasting impact connected to a gravitational wave's passage is known as gravitational-wave memory. By contrasting the two parts of a gravitational-wave signal, this difference can be comprehended. The distance between objects is momentarily stretched and compressed by the usual oscillatory signal. In contrast, the memory is a change that persists after the primary wave has passed. As a result, the impact is referred to as “permanent” or “nonoscillatory.”

Memory is expected to be much weaker than the main signal; GWM a secondary, cumulative effect from a primary energy flux. The nonlinear memory itself is not driven directly from moving masses which cause that change, but rather the energy that is carried by that gravitational wave (Heisenberg, Yunes & Zosso, 2023).

Since gravity is weak, it would cause a secondary energy flux to produce a permanent displacement smaller than the peak oscillatory amplitude in a primary wave. Issues with the instruments include: ground-based instruments such as LIGO are optimised for high frequency ripples and struggle to measure permanent, low frequency offsets. Although gravitational-wave memory is predicted theoretically, it presents a significant observational challenge. The memory component is expected to be much weaker than the primary oscillatory signal, meaning that current detectors may have difficulty distinguishing it from instrumental noise and other sources of uncertainty. As a result, research has focused not only on understanding the theoretical properties of memory but also on determining whether it can be measured with current and future gravitational-wave observatories.

Research Gap

Memory is a theoretically predicted non-oscillatory component of gravitational waves, but it has not been directly directed itself. Heisenberg, Yunes, and Zosso describe the process of gravitational-wave memory as non-oscillatory to the gravitational-wave strain; where its dominant component to the energy carried by gravitational waves and the nonlinear nature of general relativity. Several observational approaches have been investigated by researchers, and studies have determined the possibility of detecting memory through ground-based interferometers, a measurement technique that combines waves such as light, radio, or radar signals, from multiple receivers to detect structural changes, such as LIGO and Virgo, whilst the other work has investigated pulsar timing arrays. These instruments face environmental factors such as seismic activity from the earth, gravity gradients, and thermal disturbances which eventually mask the passing wave.

Detection

LIGO, VIRGO, and KAGRA are ground-based interferometric gravitational-wave detectors since each detector splits a laser beam down two long, perpendicular arms bouncing off mirrors, measuring microscopic phase shifts when a gravitational wave passes.

Gravitational-wave memory is a theoretically predicted effect of general relativity, but it has not yet been directly detected observationally. The main difficulty is that the memory component is nonoscillatory and substantially weaker than the main gravitational-wave signal (Heisenberg, Yunes & Zosso, 2023). Much of gravitational-wave astronomy has focused on detecting the oscillatory components of gravitational waves, while memory represents a different part of the signal that produces a lasting change in spacetime after the wave has passed (Grant & Nichols, 2023). The existing literature has developed several approaches for investigating whether gravitational-wave memory can be detected, including searches with ground-based interferometers, pulsar timing arrays, and proposed space-based observatories.

A major gap in the observational literature is therefore the absence of a confirmed memory detection despite increasingly sensitive gravitational-wave observations. Existing studies have instead placed limits on memory or estimated how future observations could make detection possible (Hübner, Lasky & Thrane, 2021; NANOGrav, 2020). The literature also shows that detecting memory may require combining information from many gravitational-wave events rather than relying on one typical event, because the memory signal is much weaker than the primary waveform (Lasky et al., 2016). Ground-based detectors such as LIGO, Virgo, and KAGRA can search for memory because gravitational-wave memory produces a nonoscillatory change in the detector response rather than the usual oscillating waveform. For binary black-hole mergers, the memory signal is associated with the strong-field and nonlinear behavior of general relativity, making these events important potential sources for memory searches. One challenge is that the memory signal is considerably smaller than the oscillatory gravitational-wave signal produced by the same merger, making it difficult to distinguish from detector noise and other sources of uncertainty.

In another scenario, Lasky et al. investigated whether the memory from GW150914-like binary black-hole mergers could become detectable when observations of many events were combined. Under the specific assumptions of their study, they estimated that approximately 35 binary black-hole events could produce a memory signal-to-noise ratio of about 3, while approximately 90 such events could produce a signal-to-noise ratio of about 5 at Advanced LIGO design sensitivity. These numbers are model-dependent estimates rather than a universal number of events required for detection. Combining many events can increase the statistical sensitivity to memory because the individual memory signals are weak, but this approach depends on having sufficiently accurate models of the gravitational-wave sources and their memory contributions (Lasky et al., 2016).

However, overall, searches using actual LIGO and Virgo observations have so far not produced convincing evidence for gravitational-wave memory. Hübner, Lasky, and Thrane analyzed the growing catalog of gravitational-wave observations and compared models that included memory with models that did not include memory (Hübner, Lasky & Thrane, 2021).

Other Observational approaches

Pulsar timing arrays monitor rotating neutron stars called millisecond pulsars, which beam regular radio pulses across space with stability, similar to atomic clocks. Gravitational waves stretch and squeeze across spacetime, which causes radio pulses to arrive slightly earlier or late. By comparing multiple signals, we can accurately search for a shared spatial pattern. NANOGrav, EPTA, and PPTA work under the International Pulsar Timing Array (IPTA) to combine data and increase the confidence of those measured patterns.

However, the NANOGrav 11-year dataset did not find evidence for gravitational-wave memory. Instead, the study placed upper limits on the possible strength and rate of memory events, showing that memory remains difficult to detect even when using long-term pulsar timing observations (NANOGrav, 2020). This approach is important because pulsar timing arrays observe a different frequency range from ground-based gravitational-wave detectors such as LIGO and Virgo. This means that PTAs provide a complementary method for searching for gravitational-wave memory and allow researchers to investigate memory through a different type of observation (NANOGrav, 2020).

The results from NANOGrav demonstrate that the search for gravitational-wave memory is not limited to interferometers and that different observational methods can provide different constraints on the effect. Although no memory signal has been confirmed through pulsar timing arrays, continued observations and larger datasets could improve the sensitivity of future searches (NANOGrav, 2020).

Future Based Detection

Possible gravitational-wave observations may provide new opportunities to detect gravitational-wave memory by improving detector sensitivity and observing different types of astrophysical sources. One important example is the Laser Interferometer Space Antenna (LISA), a planned space-based gravitational-wave observatory designed to detect gravitational waves in a lower-frequency range than ground-based detectors such as LIGO and Virgo. LISA will be particularly useful for observing mergers involving massive black holes, which could produce measurable memory signals (Inchauspé et al., 2025).

Unlike LIGO, Virgo, and KAGRA, which are ground-based detectors primarily sensitive to higher-frequency gravitational waves, LISA will operate in space and target the millihertz frequency band, allowing it to observe gravitational-wave signals that ground-based detectors cannot readily detect. This complementary frequency coverage will allow LISA to investigate a different range of sources, including massive black-hole mergers, extreme-mass-ratio inspirals, and compact binaries in the Milky Way (NASA, 2024; LISA Consortium, n.d.).

Inchauspé et al. (2025) investigated the possibility of measuring gravitational-wave memory with LISA by including the time-dependent response of the detector in their simulations. Their results suggest that memory could become a measurable feature in several LISA observations of massive black-hole binary mergers. This is important because it indicates that future space-based observations may provide additional opportunities to study gravitational-wave memory beyond the capabilities of current ground-based detectors.

Another possibility is the development of more sensitive ground-based detectors. Grant and Nichols (2023) examined the prospects for detecting displacement and spin memory with current and future observatories. Their study considered how improved detector sensitivity and the accumulation of observations could increase the likelihood of measuring these effects. Future detectors may therefore improve the prospects of detecting displacement memory from individual strong events, while detecting weaker effects such as spin memory may require combining information from multiple events (Grant & Nichols, 2023).

Grant and Nichols (2023) considered different sensitivity stages for the LIGO–Virgo–KAGRA network, including O4 "design" sensitivity and O5 "plus" sensitivity. These refer to different observing-run sensitivity configurations, with O5 "plus" representing a higher target sensitivity than O4 "design." Their forecasts examine how detector sensitivity and the amount of observing time can affect the prospects for detecting memory; they do not provide a detailed engineering plan specifying individual hardware upgrades. They also considered Cosmic Explorer, a proposed third-generation ground-based observatory with greater projected sensitivity (Grant & Nichols, 2023).

Spin memory differs from displacement memory in the kind of lasting effect it describes. Displacement memory refers to a lasting change in the relative separation of freely falling objects after a gravitational wave passes. Spin memory is instead associated with a cumulative relative time delay for light travelling in opposite directions around a closed path. It is related to the angular momentum carried by gravitational radiation and represents a different observable effect from the lasting displacement usually associated with gravitational-wave memory (Grant & Nichols, 2023).

However, detecting gravitational-wave memory still remains challenging because its signal is much weaker than the primary oscillatory waveform. Nevertheless, improvements in detector sensitivity and larger gravitational-wave datasets will help researchers investigate memory more effectively. A successful detection would provide an opportunity to test predictions of general relativity and examine how gravitational radiation produces lasting changes in spacetime.

Conclusion

The literature reviewed in this paper demonstrates that several approaches have been developed to investigate the possibility of detecting gravitational-wave memory. Ground-based interferometers such as LIGO and Virgo have been used to search for memory associated with binary black-hole mergers, while pulsar timing arrays provide a complementary method of investigating memory through pulsar pulse-arrival times. However, these observational studies have not yet established a confirmed detection, highlighting the continuing gap between theoretical predictions and observational evidence (Lasky et al., 2016; Hübner, Lasky & Thrane, 2021; NANOGrav, 2020).

Future observatories, particularly LISA, may improve these prospects by observing different gravitational-wave sources and providing new opportunities to measure memory. LISA is designed to observe low-frequency sources such as massive black-hole binaries, extreme-mass-ratio inspirals, and compact binaries in our Galaxy, while LIGO, Virgo, and KAGRA observe a different, higher-frequency band (NASA, 2024; LISA Consortium, n.d.). Research also suggests that combining multiple gravitational-wave events could increase the sensitivity of searches when individual signals are too weak to detect independently (Grant & Nichols, 2023; Inchauspé et al., 2025).

All the observational methods mentioned are complementary since different observatories are sensitive to different frequency changes and gravitational-wave sources. Grant and Nichols, for example, directly examine the displacement and spin memory, with displacement memory referring to a lasting change in the relative separation of initially co-moving observers and spin memory being a related effect involving an initial relative time delay for light travelling around a closed path. Ground-based detectors such as LIGO, Virgo, and KAGRA can search for memory because gravitational-wave memory produces a nonoscillatory change in the detector response rather than the usual oscillating waveform (Grant & Nichols, 2023).

Gravitational-wave memory is significant because it is connected to the nonlinear nature of general relativity and the geometry of spacetime. Measuring memory could provide an additional way to test predictions of gravity and investigate whether observations agree with general relativity or place constraints on alternative theories. Although confirmed detection remains an important goal for future gravitational-wave astronomy, continued research into its theoretical properties and observational prospects may help reveal how gravitational radiation produces lasting effects on spacetime.

Works Cited

Aggarwal, Kshitij, et al. “The NANOGrav 11 yr Data Set: Limits on Gravitational Wave Memory.” The Astrophysical Journal, vol. 889, no. 1, 2020, article 38, https://doi.org/10.3847/1538-4357/ab6083.

Chakraborty, Indranil, and Sayan Kar. “A Simple Analytic Example of the Gravitational Wave Memory Effect.” The European Physical Journal Plus, vol. 137, 2022, https://doi.org/10.1140/epjp/s13360-022-02593-y.

Flanagan, Éanna É., et al. “Persistent Gravitational Wave Observables: Nonlinear Plane Wave Spacetimes.” Physical Review D, vol. 101, no. 10, 2020, article 104033, https://doi.org/10.1103/PhysRevD.101.104033.

Grant, Alexander M., and David A. Nichols. “Outlook for Detecting the Gravitational-Wave Displacement and Spin Memory Effects with Current and Future Gravitational-Wave Detectors.” Physical Review D, vol. 107, no. 6, 2023, article 064056, https://doi.org/10.1103/PhysRevD.107.064056.

Heisenberg, Lavinia, Nicolás Yunes, and Jann Zosso. “Gravitational Wave Memory beyond General Relativity.” Physical Review D, vol. 108, no. 2, 2023, article 024010, https://doi.org/10.1103/PhysRevD.108.024010.

Inchauspé, Henri, et al. “Measuring Gravitational Wave Memory with LISA.” Physical Review D, vol. 111, no. 4, 2025, article 044044, https://doi.org/10.1103/PhysRevD.111.044044.

Johnson, Aaron D., et al. “Prospects of Detecting Nonlinear Gravitational Wave Memory.” arXiv, 2018, https://arxiv.org/abs/1810.09563.

Lasky, Paul D., et al. “Detecting Gravitational-Wave Memory with LIGO: Implications of GW150914.” Physical Review Letters, vol. 117, no. 6, 2016, article 061102, https://doi.org/10.1103/PhysRevLett.117.061102.

Talbot, Colm, et al. “Gravitational-Wave Memory: Waveforms and Phenomenology.” arXiv, 2018, https://arxiv.org/abs/1807.00990.

Zhang, P.-M., et al. “Soft Gravitons and the Memory Effect for Plane Gravitational Waves.” Physical Review D, vol. 96, no. 6, 2017, article 064013, https://doi.org/10.1103/PhysRevD.96.064013.

A Trigonometric Model of a Vertically Oscillating Arrow Passing Through a Small Hole

By Junhyeok Sung

Abstract

The passage of a flexible arrow through a fixed opening depends on the displacement of each point at its own crossing time. This paper develops a two-dimensional kinematic model in which an arrow of length L moves at constant speed v and undergoes a prescribed cosine-shaped vertical oscillation with displacement scale h_0. Substituting the crossing time into the displacement function reduces the passage condition to a maximum-value problem. When the oscillation period equals the time required to travel one arrow length, the exact radius of a centered opening is $R_{req}(D) = \frac{h_0}{2}[1 + |cos(2\pi D/L)|]$, where D is the horizontal position of the opening. The radius ranges from $h_0/2$ to h_0, with minima at odd quarter-length positions. An initial phase shifts these positions; if the phase is unrestricted, the worst-case radius is h_0 at every position. The result is conditional on the prescribed shape, synchronization, zero shaft thickness, and a thin opening plane. It establishes how spatial displacement and crossing time interact within an idealized mathematical model rather than predicting the performance of a measured arrow.

Introduction

An arrow shaft can deform during release and vibrate as it moves forward. Elasticity and harmonic motion provide a physical motivation for representing this behavior with trigonometric functions (Moebs et al.). Actual arrow motion involves more than a single vertical oscillation. Kooi and Sparenberg analyze lateral motion during interaction with a bow using moving boundary conditions (285). Rieckmann and colleagues investigate three-dimensional vibration under free–free boundary conditions using numerical models and experimental comparisons (114). These studies support treating the shaft as flexible, but they do not validate the particular cosine shape or passage formula developed here.

This paper asks how the position of a fixed, centered opening affects the minimum radius required for an idealized oscillating arrow to pass through it. The arrowhead alone cannot determine this radius: each point along the shaft reaches the opening at a different time. The relevant quantity is therefore the displacement of each point at its crossing time, rather than the shape of the entire arrow at one instant.

The analysis uses constant-speed motion, trigonometric substitution, a product-to-sum identity, and the range of the cosine function. A prescribed spatial shape makes the problem accessible without solving the equations of an elastic beam. The model first gives a general passage condition and then yields an exact formula under a specified synchronization condition. Its scope is a mathematical investigation of a chosen kinematic model; no measurements of arrow motion are used.

Model Formulation

Geometry and assumptions

The arrow is represented by a line of zero thickness in the x–y plane. Its nominal length is $L > 0$, and the coordinate $s \in [0, L]$ labels a material point measured backward from the arrowhead along the straight, undeformed shaft. Thus, $s = 0$ denotes the head and $s = L$ denotes the rear end. The arrow moves in the positive x direction with constant speed $v > 0$. Gravity, drag, damping, and rotation are omitted.

At time $t = 0$, the head has horizontal coordinate zero. The horizontal position of the point labeled s is approximated by

$$x(t, s) = vt - s. \quad (1)$$

This approximation neglects horizontal changes caused by bending. The displacement scale satisfies $h_0/L \ll 1$; for the chosen shape, the corresponding maximum slope is $2\pi h_0/L$. Equation (1) does not impose exact preservation of the length of the bent shaft.

The opening lies in an infinitesimally thin plane at $x = D$, where $D \geq 0$, and has fixed center $(D, 0, 0)$ and radius $R \geq 0$. Because the modeled motion has $z = 0$, the circular opening condition $y^2 + z^2 \leq R^2$ reduces to $|y| \leq R$. Equality permits boundary contact. A strictly positive clearance requires a larger radius than the minimum derived below.

Table 1 defines the principal variables, and Figure 1 shows the geometry.

Table 1. Variables used in the model.

Symbol	Definition	Unit
t	Time after the initial instant	s
s	Nominal distance backward from the arrowhead	m
L	Nominal arrow length	m
v	Constant forward speed	m/s
h_0	Displacement scale of the prescribed shape	m
T	Oscillation period	s
$f = 1/T$	Oscillation frequency	Hz
$\omega = 2\pi/T$	Angular frequency	rad/s
D	Horizontal position of the opening	m
R	Radius of the centered opening	m

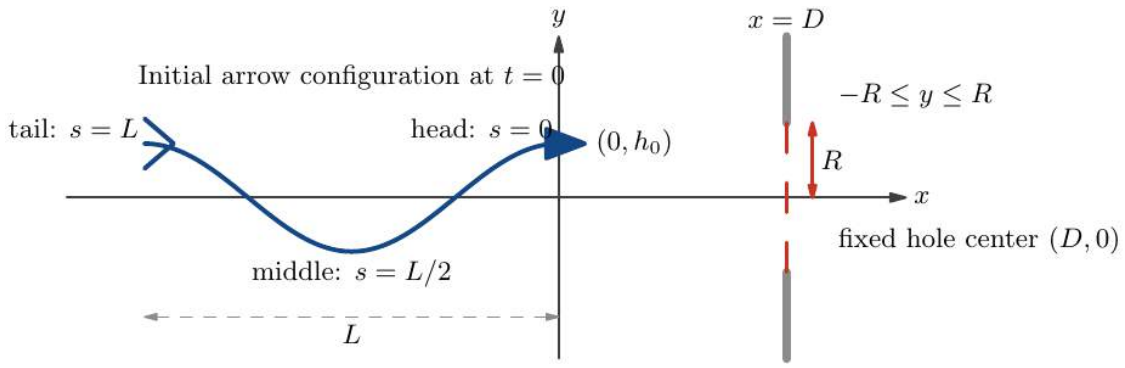


Figure 1. Coordinate system and model variables. At $t = 0$, the head is at $(0, h_0)$. In the two-dimensional view, the opening at $x = D$ extends from $y = -R$ to $y = R$.

Prescribed vertical oscillation

The spatial displacement profile is assumed to be

$$h(s) = h_0 cos\left(\frac{2\pi s}{L}\right). \quad (2)$$

It completes one cosine cycle along the shaft and gives $h(0) = h(L) = h_0$ and $h(L/2) = -h_0$. Figure 2 illustrates these values. This shape is a qualitative modeling assumption. It is not derived from material properties or the boundary conditions of a free elastic beam.

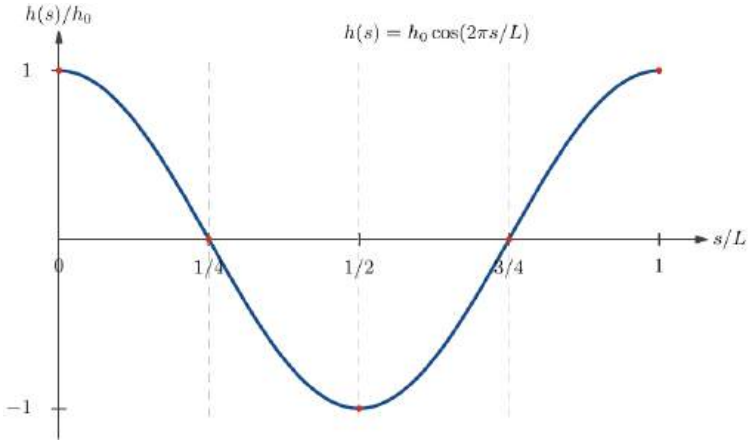


Figure 2. The assumed spatial profile in dimensionless coordinates.

Multiplying the spatial profile by a time-dependent cosine gives

$$y(t, s) = h_0 cos\left(\frac{2\pi s}{L}\right) cos(\omega t), \qquad \omega = \frac{2\pi}{T}. \quad (3)$$

The selected initial phase places the head at its maximum positive displacement at $t = 0$. At $t = T/4$, the time factor is zero and the modeled shaft is straight. At $t = T/2$, the factor is

-1 and the initial shape is vertically reversed. Figure 3 shows these configurations in coordinates attached to the shaft; forward translation is omitted from the illustration.

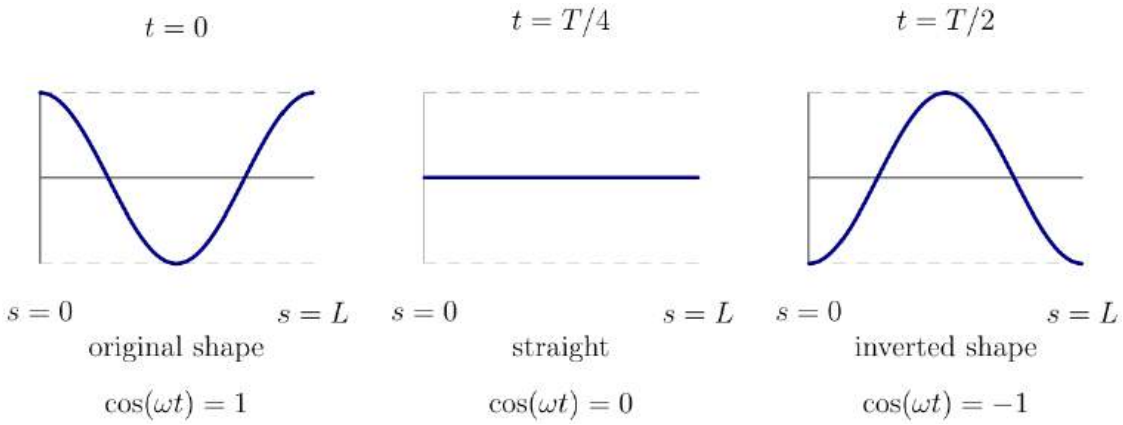


Figure 3. Spatial profiles at $t = 0$, $T/4$, and $T/2$. The time factor changes the sign and magnitude of the prescribed bending profile.

Passage Through the Opening

The point labeled s reaches the opening when $vt - s = D$. Its crossing time is therefore

$$t_D(s) = \frac{D+s}{v}. \quad (4)$$

The head crosses at D/v, whereas the rear end crosses at $(D + L)/v$. Substituting equation (4) into equation (3) gives the crossing-height function

$$y_D(s) = h_0 cos\left(\frac{2\pi s}{L}\right) cos\left(\frac{\omega(D+s)}{v}\right). \quad (5)$$

This function records a sequence of crossing events. Its values are evaluated at different times, so its graph is not an instantaneous shape of the arrow.

Every point passes through the centered opening if $|y_D(s)| \leq R$ for all $s \in [0, L]$. Since the crossing-height function is continuous on this closed interval, the required radius is

$$R_{req}(D) = \max_{0 \leq s \leq L} |y_D(s)|. \quad (6)$$

Thus, the passage problem becomes a maximum-value problem. Equation (6) applies for any positive frequency and forward speed under the stated geometric assumptions. It also gives the general bound $0 \leq R_{req}(D) \leq h_0$. A radius equal to this maximum includes all crossing points when contact with the boundary is permitted.

Exact Solution Under Synchronization

Synchronization condition

An exact expression follows when

$$\frac{\omega}{v} = \frac{2\pi}{L}, \qquad equivalently \qquad T = \frac{L}{v}. \quad (7)$$

Both sides of the first equality have units of inverse length. Under this condition, one oscillation occurs during the time required to travel forward by one arrow length. The term synchronization refers only to this equality of time scales; it does not imply resonance or an externally driven motion. It is an additional assumption, not a condition established for a real arrow.

Let $u = 2\pi s/L$ and $\beta = 2\pi D/L$. Equation (5) becomes $y_D(s) = h_0 cosucos(u + \beta)$. Applying the identity $cosAcosB = \frac{1}{2}[cos(A - B) + cos(A + B)]$ gives

$$y_D(s) = \frac{h_0}{2}[cos\beta + cos(2u + \beta)]. \quad (8)$$

Required radius

As s runs from zero to L, u runs from zero to 2π. Consequently, $2u + \beta$ spans an interval of length 4π, and its cosine attains every value in $[-1, 1]$. Writing $c = cos\beta$, the crossing heights fill the interval

$$\frac{h_0}{2}(c - 1) \le y_D(s) \le \frac{h_0}{2}(c + 1).$$

The largest distance from the fixed center $y = 0$ is therefore

$$R_{req}(D) = \frac{h_0}{2} max\{|c - 1|, |c + 1|\}.$$

For $-1 \le c \le 1$, the endpoint distances are $1 - c$ and $1 + c$, whose maximum is $1 + |c|$. Substitution yields the exact result

$$R_{req}(D) = \frac{h_0}{2}\left[1 + \left|cos\left(\frac{2\pi D}{L}\right)\right|\right]. \quad (9)$$

Since the absolute cosine lies between zero and one,

$$\frac{h_0}{2} \le R_{req}(D) \le h_0. \quad (10)$$

Both bounds are attained. The fixed center of the opening is essential to this conclusion. If its vertical position could be chosen separately for each D, the crossing interval would instead have midpoint $h_0 cos\beta/2$ and half-width $h_0/2$. An opening centered at that midpoint would have required radius $h_0/2$ for every D in the synchronized model.

Favorable and unfavorable positions

The minimum in equation (10) occurs when $cos(2\pi D/L) = 0$. For nonnegative opening positions, the solutions are

$$D = \frac{(2k+1)L}{4}, \qquad k = 0, 1, 2, ... \quad (11)$$

At these odd quarter-length positions, the crossing-height range is $[-h_0/2, h_0/2]$. The largest radius occurs at $D = \frac{kL}{2}, \qquad k = 0, 1, 2, ..., \quad (12)$

where $|cos(2\pi D/L)| = 1$ and $R_{req} = h_0$.

Figure 4 compares two crossing profiles. At $D/L = 1/4$, equation (8) gives $y_D/h_0 = -\frac{1}{2}sin(4\pi s/L)$. At $D/L = 1/2$, it gives $y_D/h_0 = -cos^2(2\pi s/L)$. The second profile remains nonpositive, but its greatest distance from the fixed center is one in normalized units. Figure 5 shows the complete radius function, whose period is $L/2$ in D.

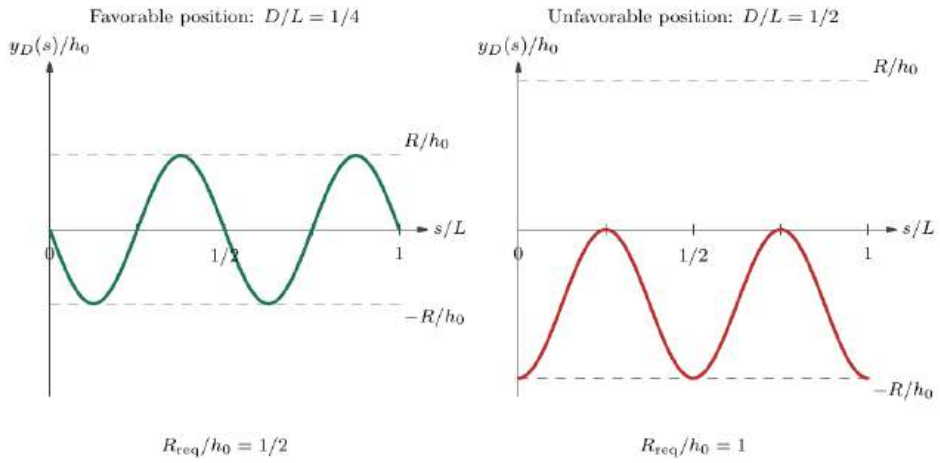


Figure 4. Crossing profiles at $D/L = 1/4$ and $D/L = 1/2$. Each point is evaluated at its own crossing time. The dashed lines show the minimum required centered-opening radius for each case.

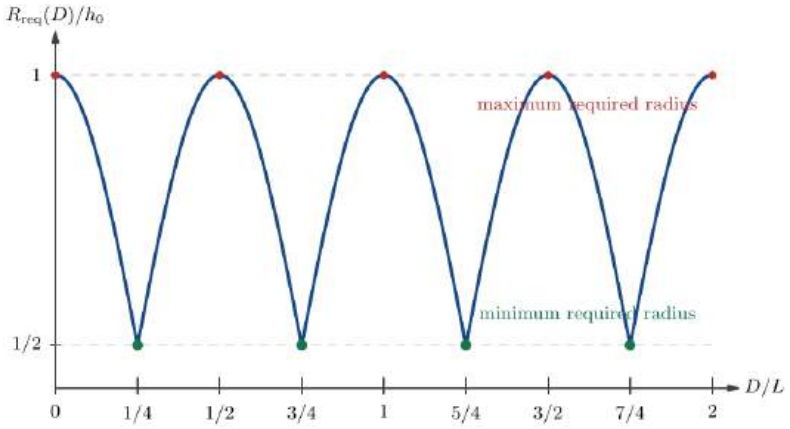


Figure 5. Normalized required radius under synchronization and zero initial phase. Minima occur at odd quarter-length positions and maxima at half-length multiples.

Numerical Illustration

Consider the illustrative parameters $L = 0.70\,m$, $v = 60\,m/s$, and $h_0 = 0.020\,m$. They are chosen to demonstrate the model and are not measurements of a particular arrow. Synchronization fixes the period, frequency, and angular frequency as

$$T = \frac{0.70}{60} \approx 0.0117\,s, \qquad f = \frac{1}{T} \approx 85.7\,Hz, \qquad \omega = 2\pi f \approx 539\,rad/s.$$

Table 2 lists the required radii from equation (9). Moving the opening from $D = L/4$ to $D = L/2$ doubles its required diameter, although the length, speed, and displacement scale remain unchanged.

Table 2. Required opening dimensions for the illustrative synchronized model.

D/L	D (m)	R_{req}/h_0	R_{req} (m)	Diameter (mm)
1/4	0.175	1/2	0.010	20
1/2	0.350	1	0.020	40
3/4	0.525	1/2	0.010	20
1	0.700	1	0.020	40

The dimensions are exact consequences of the chosen model parameters. They do not include allowances for shaft thickness, uncertainty, or a finite wall depth.

Sensitivity and Model Limitations

Initial phase

An arbitrary initial phase ϕ changes equation (3) to

$$y(t, s) = h_0 cos\left(\frac{2\pi s}{L}\right) cos(\omega t + \phi).$$

Under synchronization, the derivation of equation (9) remains valid with β replaced by $\beta + \phi$. Hence,

$$R_{req}(D, \phi) = \frac{h_0}{2}\left[1 + \left|cos\left(\frac{2\pi D}{L} + \phi\right)\right|\right]. \quad (13)$$

The favorable positions satisfy $2\pi D/L + \phi = \pi/2 + k\pi$, with integer k chosen so that $D \geq 0$. Changing the phase therefore shifts the minima along the horizontal axis. If the phase is unrestricted, a value can always be chosen that makes the absolute cosine equal to one. The worst-case radius is then

$$\max_{\phi \in [0,2\pi]} R_{req}(D, \phi) = h_0. \quad (14)$$

Thus, the radius reduction at a favorable position depends on controlling or knowing the phase. It is not guaranteed by the opening position alone.

Departure from synchronization

The dimensionless quantities $u = s/L$, $d = D/L$, and $q = \omega L/(2\pi v) = L/(vT)$ express the general model as

$$\frac{R_{req}}{h_0} = \max_{0 \leq u \leq 1} |cos(2\pi u) cos(2\pi q(d + u) + \phi)|. \quad (15)$$

Equation (15) identifies the synchronization assumption as $q = 1$. When $q \neq 1$, the equal spatial frequencies used in equation (8) no longer hold, so the radius formula and quarter-length

positions cannot be carried over without a new maximization. The general expression provides a basis for a subsequent numerical sensitivity study. The present paper does not establish error tolerances for deviations from synchronization.

Physical scope

The prescribed cosine shape and the oscillation frequency are independent inputs. A physical vibration model would determine mode shapes and frequencies from the shaft's material properties, mass distribution, and boundary conditions. Beam vibration theory discusses this dependence (Inman), while the arrow studies cited above address more detailed mechanical behavior. The assumed shape is therefore sufficient to define a kinematic problem but does not identify an actual free-vibration mode.

The approximation $x = vt - s$ also neglects changes in horizontal projection during bending. Its accuracy depends on the slope remaining small. Gravity would shift the vertical trajectory, drag would alter crossing times, and damping would reduce the time-dependent displacement. Standard treatments of damped oscillation motivate an exponentially decaying amplitude (Moebs et al.; Inman), but these effects require a revised passage calculation.
A shaft with radius $a > 0$ introduces further collision geometry. Adding a to the centerline envelope, $a + max_s|y_D(s)|$, may serve as an approximate allowance, but it is not an established sufficient clearance condition for an inclined, moving shaft. The local shaft angle, the crossing times of surface points, and the depth of a real opening must also be considered. The exact claims in equations (9)–(14) concern only the line model and an infinitesimally thin opening plane.

Conclusion

Evaluating each point at its own crossing time reduces the passage of an idealized oscillating arrow to the maximum of a crossing-height function. For the prescribed cosine shape, a fixed opening center, zero initial phase, and synchronization $T = L/v$, the required radius is given by equation (9). It varies between $h_0/2$ and h_0, with minima at odd quarter-length positions and maxima at half-length multiples.

The model explains why the maximum instantaneous displacement alone does not determine the opening radius: spatial position and oscillation phase interact during passage. The initial-phase analysis further shows that the radius reduction is conditional; an unrestricted phase requires the full radius h_0 in the worst case. These results provide an exact trigonometric solution within the stated model. Extending them to actual arrows requires a mechanically justified vibration shape, a measured relation between speed and period, and a collision model that includes shaft thickness and opening depth.

Works Cited

Inman, Daniel J. Engineering Vibration. 5th ed., Pearson, 2022.

Kooi, B. W., and J. A. Sparenberg. “On the Mechanics of the Arrow: Archer’s Paradox.” Journal of Engineering Mathematics, vol. 31, nos. 2–3, 1997, pp. 285–303. https://doi.org/10.1023/A:1004262424363.

Moebs, William, et al. University Physics Volume 1. OpenStax, 2016. https://openstax.org/details/books/university-physics-volume-1.

Rieckmann, M., et al. “Modelling the Three-Dimensional Vibration of Composite Archery Arrows under Free–Free Boundary Conditions.” Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology, vol. 226, no. 2, 2012, pp. 114–22. https://doi.org/10.1177/1754337112442273.

PM2.5 Trends in Almaty (2017–May 2026) in the Context of the CHP-2 Coal-to-Gas Conversion Project By Arailym Oral

Abstract

Air pollution, one of the most widespread problems around the globe, leads to 7 million deaths annually. The main sources of air pollution are the burning of fossil fuels, transportation emissions, and agricultural practices. This paper focuses on a hazardous air pollutant, PM2.5, that has a negative impact on the environment, human health, and therefore countries' well-being. One of the PM2.5 sources is combined heat and power plants. In the last 10 years, the CHP-2 coal-fired heat and power plant in Almaty, Kazakhstan, posed a significant hazard to the city's air quality. This led to the decision to modernize CHP-2 in 2020, followed by project implementation from 2022 to 2026, with construction beginning in 2024. This study found that average PM2.5 concentrations declined between 2021 and 2025, falling by approximately 2.5 times from 40.99 μg/m³ to 16.87 μg/m³. However, PM2.5 concentrations increased again in 2026. These findings suggest that the period of CHP-2 modernization may be associated with improvements in air quality in Almaty, although other environmental and meteorological factors may have contributed to the observed decline. The findings can help evaluate air-quality policies and support future decisions on the modernization of major emission sources in Almaty and other cities facing similar air-pollution challenges.

Keywords: Environmental Science, Air Pollution, PM2.5, Almaty, CHP-2

Introduction

Fine particulate matter (PM2.5) refers to airborne particles with a diameter of 2.5 micrometers or smaller (World Health Organization). PM2.5 is among the most extensively studied air pollutants and has a number of negative effects; it can lead to respiratory and cardiovascular diseases, lung cancer, and strokes, along with the degradation of infrastructure through corrosion (Ness et al.). Globally, in the period between 2000 and 2019, about one million annual premature deaths were attributed to short-term exposure to particulate matter (PM2.5) (Muratuly et al.). Moreover, it poses a heavy burden on the economy. In Almaty city, a yearly pecuniary loss of US$486 million due to healthcare costs was attributed to PM10 and PM2.5 pollution (World Bank).

In comparison to other cities in Kazakhstan, Almaty usually experiences higher average pollution concentrations (Bekbossinova and Niyazbekov). Almaty is located in the foothills of the northern slope of the Trans-Ili Alatau ridge of the Tien Shan mountain system. The topography and geography of the area provide conditions for inversion layers to form, which can contribute to daily air pollution by restricting the vertical dispersion of pollutants (Kerimray et al.). Furthermore, Almaty’s population is growing, leading to higher energy use and carbon emissions (Bekbossinova and Niyazbekov).

Large combustion sources are important contributors to PM2.5 exposure in Almaty, particularly the city's combined heat and power plants (Zakarin et al.). Almaty's CHPPs generate over 80% of the city's power (Karatayev and Clarke). According to the World Bank, these sources account for approximately 25 μg/m³, or 40% of annual mean population exposure to PM2.5 in Almaty, with most of this contribution attributed to the old coal-fired CHP-2 plant, which lacked adequate emission-control equipment (World Bank).

In this regard, the modernization of CHP-2 has become a major environmental intervention in Almaty. In 2020, the Government of Kazakhstan approved a project to convert CHP-2 from coal to gas (Government of the Republic of Kazakhstan). The active phase of work began in 2024, while the completion of the project and the launch of new gas capacity are planned for 2026 (Samruk-Energy). At the time of this study, the modernization was not fully completed, while construction of the new gas station was still underway. Thus, the research period 2017-2026 does not include the full period before and after the intervention, but both the period before modernization and the period of its implementation.

Previous studies have researched the seasonal and meteorological characteristics of PM2.5 pollution in Almaty, as well as temporal variation and the influence of air-mass transport (Tursumbayeva et al.; Mukhtarov et al.). However, long-term trends in PM2.5 concentrations in Almaty, particularly in relation to major environmental policy interventions, have been less explored. Therefore, this paper examines how PM2.5 concentrations changed in Almaty from 2017 to May 2026 and whether these trends correspond with the implementation timeline of the CHP-2 coal-to-gas conversion project.

Methods

This research investigated open data from AirDataKz to analyze PM2.5 concentration trends for the period from 2017 to 2026. The data were assessed in advance to identify gaps and erroneous values. Before the analysis, the data were checked for missing PM2.5 values and errors in the data format. Periods without available measurements were treated as missing data and not as zero values. Extremely high PM2.5 values were not removed unless they were identified as errors.

The data were obtained from AirData.kz, an open air-quality data platform that combines measurements from government monitoring stations, low-cost sensors and international data aggregators. Its sources include KazHydroMet, AirGradient, OpenAQ, WAQI and the historical AirKaz network (AirData.kz, "The Data").

The availability of PM2.5 data varied throughout the study period. In 2017, data were available from 16 stations, with 2,778 hourly measurements. In 2018 and 2019, 22 and 19 stations were available, with 5,616 and 4,699 hourly measurements, respectively. In 2020, the dataset included 42 stations and 46,986 hourly measurements. From 2021 to 2025, the number of stations ranged from 12 to 16, with 74,219 to 95,584 hourly measurements recorded annually. In 2026, data from 153 stations and 149,912 hourly measurements were available.

The seasonal analysis and graphs were grouped by season: winter (December to February), spring (March to May), summer (June to August), and autumn (September to November) (National Centers for Environmental Information). Annual and seasonal average PM2.5 concentrations were calculated directly from the available hourly measurements, separately for each year, season and geographic cluster. Periods without available measurements were not included as zero values in the calculations. The annual and seasonal average PM2.5 concentrations were then visualized using graphs and compared with the timeline of the gasification of CHP-2.

In addition, this paper analyzes the average PM2.5 concentrations for different cluster areas in Almaty, utilizing a heatmap. Consequently, this research aims to examine changes in PM2.5 concentrations in Almaty over the last ten years and explore whether these changes correspond with recent CHP-2 policy interventions. However, the observed changes in PM2.5 concentrations cannot be directly attributed to the CHP-2 conversion project based on this study alone.

Results

The analysis revealed that 2018 had the highest average annual PM2.5 concentration, at 52.33 µg/m³. 2017 and 2019 had approximately similar values, and PM2.5 concentration fell by almost 10 to 41.83. From 2019 to 2020, PM2.5 concentration remained relatively constant, followed by a decline in 2022. From 2023 to 2024, it showed a gradual and minor decrease. The 2026 data involved measurements for the period from January to May, and it recorded a sharp increase in PM2.5 concentrations to 31.24 µg/m³ compared with 16.87 µg/m³ in 2025. The largest decrease was observed from 2021 to 2025 by about 2.5 times, from 40.99 to 16.87.

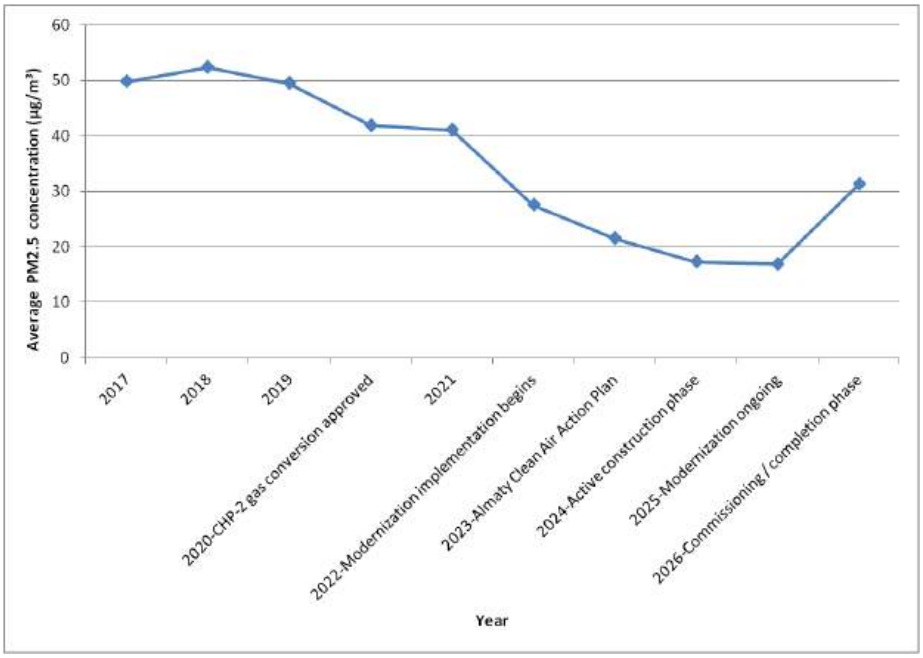


Figure 1. Annual Mean PM2.5 Concentrations in Almaty, 2017–2026, with the timeline of key CHP-2 policy interventions. The graph was generated using Microsoft Excel from AirDataKz data. Annual mean PM2.5 concentrations were calculated as the mean of the available observations for each year.

Figure 2 illustrates the annual seasonal change in PM2.5 concentration throughout the given period. Summer, autumn, and winter PM2.5 average measurements are absent for 2026 as the data is not available. Throughout the whole period, winter readings were considerably higher than those of other seasons. Summer, on the other hand, indicated the lowest PM2.5 concentrations throughout the years in comparison to other seasons. Summer $PM_{2.5}$ observations gradually decreased from 21.73 μg/m³ in 2017 to 7.92 μg/m³ in 2025, resulting in about a 64% decline. Winter and autumn reached their peaks in 2018, at 119.52 μg/m³ and 58.24 μg/m³, respectively. Winter measurements reached their minimum in 2024, at 29.34 μg/m³, which was about four times lower than the maximum level reached in 2018. Spring peaked in 2019 at about 34.47 μg/m³ and almost reached this level again in 2026, at about 30.06 μg/m³.

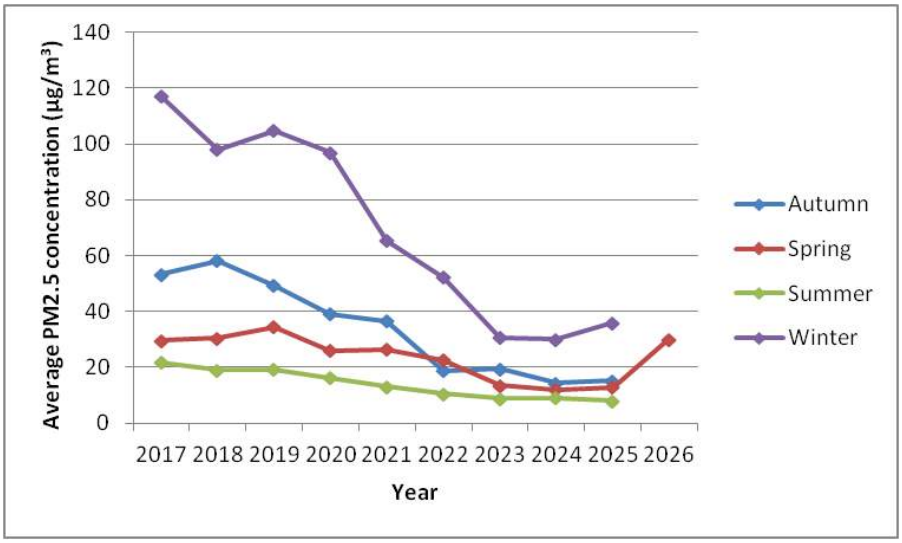


Figure 2. Seasonal Mean PM2.5 Concentrations in Almaty, 2017–2026. The line graph was created by the author using Microsoft Excel based on AirDataKz data. Seasonal mean PM2.5 concentrations were calculated using the available observations for autumn, spring, summer, and winter for each year from 2017 to 2026.

The table 1 indicates substantial spatial heterogeneity in PM2.5 concentrations. It illustrates that the earlier years of the study period experienced the highest concentrations of PM2.5 particles. The Far North cluster showed the highest PM2.5 concentration of 124.87 μg/m³ in 2017, and an observed decrease to 42.48 μg/m³ in 2026. West Central had comparably high PM2.5 concentrations in 2017–2019, 79.99 μg/m³, 77.28 μg/m³, and 94.23 μg/m³, respectively. It reached its peak in 2019, then decreased by more than half in 2020 to 42.99 μg/m³. Throughout the next several years, it declined, dropping to 3.70 μg/m³ and climbing again to 38.62 μg/m³ in 2026. West Residential had fluctuating values from 2017 to 2020. Nevertheless, it decreased by almost a half to 27.21 μg/m³ in 2025 and remained steady in 2026. Overall, several clusters showed lower concentrations in the later observed years compared with earlier years.

However, comparisons between clusters across the entire study period are limited by differences in data availability. Some clusters had missing observations for multiple years. For example, West Residential and South Central lacked data for the period between 2021–2023, while Far North had no data between 2021 and 2025, and Far West and East Suburbs had observations only for 2026. Therefore, the heatmap does not reflect continuous measurements for all clusters across the study period

Cluster \ Year	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026
City Center	44,39	51,02	47,83	43,21	42,54	38,61	37,84	30,81	22,71	25,59
East Suburbs (Talgar)										32,77
Far North	124,87	72,81	56,14	70,16						42,48
Far West										37,66
North Central	60,76	66,65	64,73	42,22	41,67	17,48	8,76	9,35	19,38	42,21
South Central	47,29	40,62	44,74	41,62					26,59	23,86
South Foothills	27,53	32,34	30,62	26,77	16,58	13,99	12,62	10,42	10,18	18,70
Southwest Suburbs	64,41	45,39	42,14	50,17						25,88
West Central	79,99	77,29	94,22	42,99	48,62	39,44	29,13	13,76	3,70	38,62
West Residential	51,09	53,00	42,65	48,79					27,22	27,08

Table 1. Spatial and Temporal Variation in PM2.5 (μg/m³) Concentrations Across Almaty, 2017–2026. The table was created by the author using Microsoft Excel based on AirDataKz data.

Discussion

The analysis indicates that the PM2.5 values showed an overall declining trend throughout the observed timeline, with substantial seasonal and spatial changes remaining. The reduction in PM2.5 levels coincided with the implementation of the CHP-2 modernization project, as the government approved the conversion of CHP-2 from coal to gas in 2020; project implementation began in 2022, with the active construction phase taking place in 2024. The decline in annual PM2.5 concentrations, particularly after 2021, therefore overlaps with the period of CHP-2 modernization, but it cannot be certainly claimed that the reason for the decline was the CHP-2 modernization, due to other factors such as transport, other emissions, and meteorological conditions.

Analysis of the spatial distribution showed that PM2.5 is unevenly distributed across Almaty. Some clusters exhibited higher average concentrations than others and lacked observations for particular years. It is particularly significant that the Far North cluster exhibited the highest observed value in 2017, even though it does not correspond to the zone of direct impact from CHP-2. It demonstrates that the spatial dynamics of PM2.5 are complex and cannot be attributed solely to emissions from CHP-2, and might be high because of different types of emissions and factors. In addition, West Central, which is located closest to CHP-2, recorded high PM2.5 concentrations in 2017–2019. The concentration decreased in 2020 and then declined further by about two times in 2023–2024, during the construction phase of the new gas-fired facility. However, since the conversion from coal to gas had not been completed during the study period, this decline cannot be considered evidence of the effectiveness of the CHP-2

modernization. Other factors, such as emissions from other sources and meteorological conditions, might also have contributed to the observed changes.

Seasonal findings revealed that the highest average PM2.5 concentration was recorded in winter. This finding is consistent with other studies on Almaty, which also indicated that winter concentrations were higher (Kerimray et al.; Tursumbayeva et al.). One possible reason might be the increased emissions from heating during colder months. In winter, temperature inversions can also prevent pollutants from dispersing (Tursumbayeva et al.). Moreover, summer recorded the lowest concentration of particles, which might be related to lower heating demand and better conditions for pollutant dispersion.

Conclusion

Overall, this study showed a substantial decrease in PM2.5 concentrations in Almaty, especially in the period from 2021 to 2025, along with obvious seasonal and spatial differences. The decrease coincided with the modernization of CHP-2. However, the available data do not allow us to determine the individual effect of the modernization of CHP-2, since levels of PM2.5 levels may also be affected by other sources of emissions and meteorological conditions. Further monitoring is needed after the modernization is fully completed in order to better assess its impact on air quality in Almaty.

Future research could compare PM2.5 concentrations before and after the completion of CHP-2 modernization while accounting for seasonal and meteorological changes. Other sources of emissions, such as traffic and residential heating, could also be studied to better understand the factors contributing to air pollution in Almaty.

Limitations

This study has potential limitations. One limitation of this study is that the modernization of CHP-2 was not fully completed during the study period. Therefore, the 2017–2026 study period does not include a complete before-and-after period for the intervention, but instead covers the period before the modernization and the period during its implementation. This means that the full effect of the modernization on PM2.5 concentrations cannot yet be evaluated.

Another limitation is that the analysis for 2026 is based only on the data available for January through May, meaning that the 2026 results should not be directly compared with the full-year results of previous years.

A further limitation is the difference in data availability between clusters. Some clusters had missing observations for multiple years, which limits comparisons across the entire study period. The heatmap does not represent continuous measurements for all clusters throughout the study period, and comparisons between clusters should be interpreted with caution.

The number of monitoring stations and available measurements varied across the study period. Therefore, changes in annual mean PM2.5 concentrations should be interpreted with caution, as they may partly reflect differences in data availability and monitoring coverage.

Works Cited

AirData.kz. “The Data.” AirData.kz, https://airdata.kz/en/data/. Accessed 14 Oct. 2026.

AlES JSC. “Modernization of Almaty CHPP-2.” AlES, https://www.ales.kz/en/modernization-of-almaty-chpp-2/. Accessed 13 Oct. 2026.

AlES JSC. “Modernization of Almaty CHPP-2 with Minimized Environmental Impact.” AlES, https://www.ales.kz/en/information-on-the-approval-of-the-category-and-the-implementation-of-the-pqs-2-0-procedure/. Accessed 13 Oct. 2026.

Bekbossinova, A., and A. Niyazbekov. “The Impact of Urbanization on Air Quality in Largest Cities of Kazakhstan.” Eurasian Journal of Economic and Business Studies, vol. 68, no. 3, 2024, pp. 66–81, https://doi.org/10.47703/ejebs.v68i3.419.

Government of the Republic of Kazakhstan. “Government Approved Project of Gas Supply to Almaty CHPP-2.” Prime Minister of the Republic of Kazakhstan, 20 Oct. 2020, https://primeminister.kz/en/news/v-pravitelstve-odobren-proekt-gazifikacii-almatinskoy-tec-2-2094726.

Karatayev, M., and M. L. Clarke. “A Review of Current Energy Systems and Green Energy Potential in Kazakhstan.” Renewable and Sustainable Energy Reviews, vol. 55, 2016, pp. 491–504, https://doi.org/10.1016/j.rser.2015.10.078.

Kerimray, Aiymgul, et al. “Spatiotemporal Variations and Contributing Factors of Air Pollutants in Almaty, Kazakhstan.” Aerosol and Air Quality Research, vol. 20, no. 6, 2020, pp. 1340–52, https://doi.org/10.4209/aaqr.2019.09.0464.

Mukhtarov, R., et al. “An Episode-Based Assessment for the Adverse Effects of Air Mass Trajectories on PM2.5 Levels in Astana and Almaty, Kazakhstan.” Urban Climate, vol. 49, 2023, article 101541, https://doi.org/10.1016/j.uclim.2023.101541.

Muratuly, A., et al. “Urban PM2.5 Pollution in Kazakhstan: Health Burden and Economic Costs.” Environmental Science: Advances, vol. 5, no. 1, 2026, pp. 281–91, https://doi.org/10.1039/D5VA00194C.

National Centers for Environmental Information. “Meteorological versus Astronomical Seasons.” National Centers for Environmental Information, 22 Sept. 2016, https://www.ncei.noaa.gov/news/meteorological-versus-astronomical-seasons.

Ness, J. E., G. Heath, and V. Ravi. An Overview of Policies Influencing Air Pollution from the Electricity Sector in Central Asia. National Renewable Energy Laboratory, 2022, NREL/TP-7A40-81861, https://www.nrel.gov/docs/fy23osti/81861.pdf.

Samruk-Energy JSC. Integrated Annual Report 2024. Samruk-Energy, 2024, https://ar2024.samruk-energy.kz/en/.

Tursumbayeva, Madina, et al. “Planetary Boundary Layer and Its Relationship with PM2.5 Concentrations in Almaty, Kazakhstan.” Aerosol and Air Quality Research, vol. 22, no. 8, 2022, article 210294, https://doi.org/10.4209/aaqr.210294.

World Bank. Clean Air and Cool Planet: Volume II: Integrated Air Quality Management and Greenhouse Gas Reduction for Almaty and Nur-Sultan. 2022,

https://documents1.worldbank.org/curated/en/099920008292227204/pdf/P1708700f4b6f30f0bf1a05fe6c088bdd2.pdf.

World Bank. Towards Cleaner Industry and Improved Air Quality Monitoring in Kazakhstan. 2013, https://documents.worldbank.org/en/publication/documents-reports/documentdetail/304541468292529914.

World Health Organization. “WHO Global Air Quality Guidelines.” World Health Organization, 22 Sept. 2021, https://www.who.int/news-room/questions-and-answers/item/who-global-air-quality-guidelines.

Zakarin, E. A., et al. “Simulation of Air Pollution in Almaty City under Adverse Weather Conditions.” Russian Meteorology and Hydrology, vol. 46, no. 2, 2021, pp. 121–28, https://doi.org/10.3103/S1068373921020072.

Why Are Three-Point Shots Mathematically Advantageous in the NBA? By Taehoon Kim

An Analysis of Modern Basketball Offensive Strategy Using Expected Value

Abstract

This paper investigates why three-point shots have become such an important part of modern NBA offense from a mathematical point of view. I mainly use expected value to compare the efficiency of two-point and three-point shots. Even though three-point shots are more difficult to make, their higher point value can make them more efficient when the shooting percentage is high enough. I also examine actual NBA games involving the Boston Celtics, Golden State Warriors, and Houston Rockets to see how this idea works in real situations. These examples show that both shooting percentage and the number of attempts can affect a team's scoring advantage. However, the Houston Rockets' 27 consecutive missed three-pointers also show that a strategy with a good, expected value can still fail in a single game because of variability. I then consider statistics such as effective field goal percentage, true shooting percentage, and the Four Factors to show that real basketball is more complicated than a simple expected-value calculation. Overall, this paper shows how probability and expected value can help explain the NBA's increasing use of three-point shots.

Keyword Expected Value, Three-Point Shot, NBA, Probability, Shooting Efficiency, eFG%, Basketball Analytics

1. Introduction

The way NBA teams play offense has changed a lot over the past two decades. In the past, teams relied more on post plays, shots near the basket, and mid-range jump shots. This made sense because shots closer to the basket were generally easier to make, so many teams considered two-point shots a safer and more reliable choice.

But when I looked at how NBA basketball has changed, one thing stood out to me immediately: teams now take far more three-point shots. Data analysis has also become a much bigger part of basketball, causing teams to think differently about what actually makes a shot efficient.

Stephen Curry and the Golden State Warriors are probably the clearest example of this change. Curry regularly takes three-point shots from distances that would have seemed like bad shot choices in the past, but he has been able to make them at a high rate. The Warriors also won multiple championships while using an offense that depended heavily on three-point shooting. Their success helped other teams recognize the value of the three-point shot.

At first, I thought the explanation was mostly that Curry is simply an unusually good shooter. But there is also a basic mathematical reason why three-point shooting can be so effective. The idea can be explained using expected value. This led me to the following research question: Can a high volume of three-point attempts improve an NBA team's offensive efficiency

and chance of winning from the perspective of expected value, even though three-point shots have a lower probability of being made?

To explore this question, I first compare the expected values of two-point and three-point shots. Then I look at several actual NBA games where three-point shooting had a noticeable effect on the result. I also include an example, where relying heavily on three-point shots went very badly. This is important because expected value describes what should happen on average, not what must happen in every game.

2. Theoretical Background

2.1 Expected Value and Basketball

Expected value is a way of describing the average result we would expect if the same situation happened many times.

For a basketball shot, I can use the simple formula:

$$\mathbf{EV = P \times V}$$

where:

- EV = expected points from one shot
- P = probability of making the shot
- V = number of points received if the shot is made

For example, suppose a player makes 50% of his two-point shots.

$$\mathbf{EV_2 = 0.50 \times 2 = 1.00}$$

So, over many similar attempts, that player would average about 1 point for each shot.
Now consider a player who makes 33.3% of his three-point shots.

$$\mathbf{EV_3 = 0.333 \times 3 \approx 0.999}$$

This is also almost exactly 1 point per attempt.

I found this comparison interesting because the difference in shooting percentage seems very large at first. One player makes half of his shots, while the other makes only about one-third. But their expected points are almost the same because a made three-pointer is worth more. We can also find the point where a two-point shot and a three-point shot have the same expected value. Let p_2 be the probability of making a two-point shot and p_3 be the probability of making a three-point shot.

The expected values are equal when:

$$\mathbf{2p_2 = 3p_3}$$

Solving for p_3:

$$\mathbf{p_3 = (2/3)p_2}$$

If a player makes 50% of his two-point shots:

$$\mathbf{p_3 = (2/3)(0.50)}$$
$$\mathbf{p_3 \approx 0.3333}$$

So, a 50% two-point shot and a 33.3% three-point shot have almost the same expected value.

This also means that once the three-point percentage goes above 33.3%, the three-pointer can become more valuable than a 50% two-point shot.

For example, suppose a player shoots 36% from three-point range.

$$\mathbf{EV_3 = 0.36 \times 3 = 1.08}$$

A 50% two-point shooter still has:

$$\mathbf{EV_2 = 0.50 \times 2 = 1.00}$$

The difference is only 0.08 points per shot. When I first calculated this, a difference of only 0.08 points did not seem important at all. But if the difference continues over 100 attempts:

$$\mathbf{100(1.08) - 100(1.00) = 8}$$

the difference becomes 8 expected points.

Eight points can make a significant difference in an NBA game. This helped me understand why a very small difference in efficiency per shot can become essential when teams repeat the same types of decisions many times.

3. Analysis of Actual NBA Games

3.1 The Boston Celtics: The Effect of High Three-Point Volume

One game I looked at was the Boston Celtics against the Los Angeles Clippers on November 25, 2024.

Three-point shooting played a major role in Boston's large victory. The Celtics attempted far more three-point shots than the Clippers and finished with a **45-point advantage in scoring from three-point shots**.

At first, it is easy to focus only on a team's shooting percentage. But the number of attempts is also important. Suppose a team takes n three-point shots and has a probability p of making each one.

The expected number of points from those shots can be written as:

$$E = 3np$$

This shows that if the shooting percentage stays reasonably similar, taking more three-point shots can increase a team's expected total points.

For example, suppose two teams both shoot 35% from three-point range.

If Team A takes 20 three-point shots:

$$E_A = 20 \times 0.35 \times 3 = 21$$

The expected three-point scoring total is 21 points.

Now suppose Team B takes 40 three-point shots at the same percentage:

$$E_B = 40 \times 0.35 \times 3 = 42$$

Its expected total becomes 42 points.

Of course, this does not mean Team B will score exactly 42 points in one real game. It could score much more or much less. The calculation only describes an average result if the same situation is repeated many times.

The Celtics game clearly illustrates why volume can matter. A team that creates many more three-point opportunities can build a large scoring advantage if it can also maintain a reasonable percentage.

3.2 Golden State Warriors: Can Volume Matter More Than Percentage?

Another game that interested me was the Golden State Warriors against the Phoenix Suns on February 6, 2026.

The Warriors played without Stephen Curry. Because Curry is such an important three-point shooter, I expected that Golden State might take fewer three-pointers or struggle more from outside.

Instead, the Warriors continued to take a very large number of three-point shots. According to the game statistics I examined, Golden State attempted **55 three-pointers and made 36.4% of them**. Phoenix attempted **43 three-pointers and made 37.2%**.
Phoenix actually had the better shooting percentage.

However, when I calculated the approximate number of made three-pointers, the result told a different story.

For Golden State:

55 × 0.364 ≈ 20.02

This is approximately 20 made three-pointers.

20 × 3 = 60 points

For Phoenix:

43 × 0.372 ≈ 16.00

This is approximately 16 made three-pointers.

16 × 3 = 48 points

Golden State therefore scored approximately 12 more points from three-pointers even though Phoenix had the slightly higher percentage.

This changed the way I thought about shooting efficiency. I originally assumed that the team with the higher three-point percentage would usually have the advantage. However, percentage alone does not tell the whole story.

The difference in shooting percentage was only:

37.2% − 36.4% = 0.8 percentage points

The difference in attempts was:
55 − 43 = 12 attempts

Golden State created 12 additional opportunities to score from three-point range. In this game, that additional volume mattered more than the small difference in percentage.
The Warriors eventually won 101–97. The final margin was only four points, so their advantage in three-point scoring seems especially important when looking at the final result.

4. Does an Expected-Value Strategy Always Work?

4.1 The 2018 Houston Rockets and 27 Consecutive Missed Three-Pointers

So far, the examples make a three-point-heavy strategy look very strong. But expected value does not mean that a strategy will work every time.

One of the clearest examples of this limitation was Game 7 of the 2018 Western Conference Finals between the Houston Rockets and the Golden State Warriors. Houston had won 65 games during the regular season and was one of the NBA teams most committed to shooting large numbers of three-pointers. In one of the most important games of the season, however, the strategy produced a very different result.

The Rockets attempted 44 three-point shots and made only 7.

Their three-point percentage was:

(7/44) × 100 ≈ 15.9%

Their actual average scoring from each three-point attempt was:

(7 × 3)/44 = 21/44 ≈ 0.477 points per attempt

This was an extremely low scoring rate.

The most shocking part was that Houston missed **27 consecutive three-point attempts** during the game.

This example helped me recognize an important limitation of looking only at expected value. Expected value is a long-run average. It does not tell us exactly what will happen in one game.

For example, suppose a three-point shot has a 36% chance of going in.

0.36 × 3 = 1.08 points

The expected value is 1.08 points.

But a player never actually scores 1.08 points on one shot. He normally scores either 0 or 3.That difference is important. In a small number of attempts, actual results can be very far from the expected value.

So, Houston's strategy was not necessarily mathematically wrong just because it failed in this game. If a team repeatedly chooses shots with higher expected values over hundreds of games and thousands of possessions, the strategy can still make sense.
The problem is that Game 7 is not thousands of possessions. It is one game. A strategy can be better in the long run and still fail at the worst possible time.

5. Expected Value and Variability

After looking at the Houston game, I started to think that expected value by itself was not enough to describe the whole problem.
Consider two imaginary shots.

Shot A
A 50% chance of making a two-point shot:

$$\mathbf{EV_A = 0.50 \times 2 = 1.00}$$

Shot B

A 33.3% chance of making a three-point shot:

$$\mathbf{EV_B = 0.333 \times 3 \approx 1.00}$$

The expected values are almost identical.

However, the actual possible outcomes of each shot are different.
Shot A gives either 0 or 2 points. Shot B gives either 0 or 3 points.
So even though their averages are about the same, the three-point shot can create a larger change in the score from a single attempt.

This variability can work both ways. If a team makes several three-pointers in a short period of time, the score can change very quickly. But if the team misses several in a row, it can also go through a long stretch without scoring.

The Celtics and Rockets examples demonstrate these two sides of three-point variance. Boston's high volume of three-point shooting helped create a huge scoring advantage. Houston depended heavily on the same kind of shot, but when those shots stopped falling, the result was completely different.

Because of this, I do not think it is correct to say: **Higher expected value = guaranteed victory**. It makes more sense to say: **Higher expected value = a strategy that should produce a higher average scoring result when it is repeated many times.** That difference between an average and a guarantee is one of the most important things I learned while working on this topic.

6. Going Beyond Simple Expected Value

6.1 Effective Field Goal Percentage (eFG%)

Expected value gives a simple way to compare two-point and three-point shots, but actual basketball is more complicated.
One statistic that tries to account for the extra value of a three-pointer is **Effective Field Goal Percentage (eFG%)**.

Normal field-goal percentage counts a made two-pointer and a made three-pointer in exactly the same way. Both count as one made field goal.
But they obviously do not have the same scoring value. One is worth 2 points and the other is worth 3.

eFG% adjusts for this difference:

eFG% = [FGM + 0.5(3PM)] / FGA

where:

- FGM = total field goals made
- 3PM = three-point field goals made
- FGA = total field-goal attempts

For example, suppose two players each take 10 shots.

Player A makes five two-point shots. His normal field-goal percentage is 50%, and he scores 10 points. Player B makes four three-point shots. His field-goal percentage is only 40%, but he scores 12 points.

If I only looked at normal field-goal percentage, I would say Player A shot better. But Player B actually scored more points with the same number of attempts.

For Player B:

eFG% = [4 + 0.5(4)] / 10
eFG% = 6/10 = 60%

This is why normal field-goal percentage can sometimes be misleading. Player B made a smaller percentage of his shots but produced more points from the same number of attempts.

6.2 True Shooting Percentage (TS%)

Another statistic that goes one step further is **True Shooting Percentage (TS%)**. Players do not score only from field goals. They can also score from free throws. Because of that, if I want to compare overall scoring efficiency, it makes sense to include free throws too.

True Shooting Percentage considers two-point shots, three-point shots, and free throws when measuring scoring efficiency.

This gives a broader picture than simply asking whether a player shoots a lot of three-pointers. It looks at how efficiently the player turns his scoring opportunities into actual points.

This also made me realize that the simple equation I started with,

EV = P × V

is useful, but it is only a starting point. Real basketball has more variables than one shot probability and one point value.

7. The Four Factors and Actual Basketball Games

Basketball analyst Dean Oliver introduced a framework called the **Four Factors**. It looks at four major parts of basketball efficiency:

1. Shooting efficiency
2. Turnovers
3. Offensive rebounding
4. Free-throw rate

Looking at these factors helped me understand why simply taking as many three-point shots as possible is not always the best strategy.

For example, a team might be very efficient when it actually gets a three-point attempt, but that does not help much if the team keeps turning the ball over before taking a shot.

Offensive rebounds matter for a different reason. If a team misses a shot but gets the rebound, it receives another chance to score. So, one missed shot does not always mean the possession is over.

Free throws also matter. A good free-throw shooter can score efficiently without facing normal defensive pressure.

Because of these factors, the expected points from an offensive possession cannot always be described only as:

Expected points = P(shot made) × points

Whether the team keeps possession, gets another shot, turns the ball over, or reaches the free-throw line also matters.

For me, this was an important extension of the original expected-value concept. The simple model explains why three-point shots can be valuable, but it does not explain everything that happens during an actual basketball possession.

8. Analysis: What Is the Real Advantage of the Three-Point Shot?

After comparing these examples, I do not think the main mathematical advantage of the three-point shot comes simply from the fact that 3 is greater than 2.
The more important idea is the relationship between the **probability of success and the reward for success**.

A three-point shot is harder to make than many two-point shots, but it gives 1.5 times as many points when it goes in. If the shooting percentage is high enough, the extra point can make up for the lower probability.

This becomes especially interesting when three-pointers are compared with mid-range two-point shots.
Suppose a player makes 45% of his mid-range two-point shots.

EV_mid = 0.45 × 2 = 0.90 points per attempt

Now suppose the same player makes 35% of his three-point shots.

EV_3 = 0.35 × 3 = 1.05 points per attempt

The mid-range shooting percentage is 10 percentage points higher, which might initially make it seem like the better shot.

But the expected-value calculation says something different.

1.05 − 0.90 = 0.15 expected points per attempt

Over 100 similar attempts:

100 × 1.05 = 105 expected points

while:
100 × 0.90 = 90 expected points

The difference is:
105 − 90 = 15 expected points

I think this calculation gives one of the clearest explanations for why long two-point shots have become less common in the NBA while three-point attempts have increased. A player does not have to shoot three-pointers at the same percentage as two-pointers for them to be worthwhile.

Still, the Houston Rockets example prevents me from concluding that teams should simply take as many three-pointers as possible. A higher expected value does not guarantee a better result in one game.

A better goal might be: **Create as many high-expected-value scoring opportunities as possible while recognizing that the actual result of a single game can still vary greatly.**

In the end, the NBA's three-point revolution cannot be explained only as a trend or only by the success of Stephen Curry. Curry and the Golden State Warriors showed how powerful this style of basketball could be, but expected value gives a mathematical explanation for why the strategy makes sense over many possessions.

At the same time, the 2018 Houston Rockets show the other side of probability. Even a mathematically reasonable strategy can produce a terrible result in one game.

This is what I found most interesting about the mathematics of basketball. Math cannot tell us for certain which team will win a game. What it can do is help explain why one decision may be better than another when the same type of decision is made again and again.

9. Conclusion

In this paper, I tried to understand why three-point shots have become such an important part of NBA basketball using expected value and probability. Before researching this topic, I thought the increase in three-point shooting was mainly because there are greater three-point shooters today. However, after doing the calculations, I found that there is also a mathematical reason behind this change.

The most interesting result for me was that a three-point shot does not need the same shooting percentage as a two-point shot to produce the same expected points. For example, a 33.3% three-point shot has about the same expected value as a 50% two-point shot. I also found that a 35% three-point shot gives an expected value of 1.05 points, while a 45% two-point shot gives only 0.90 points. At first, this difference did not seem very large to me, but when many shots are taken during a game or over a season, the difference can become much more important.

The actual NBA games also helped me understand that the number of three-point attempts matters. A team does not always need to have a better three-point percentage to gain an advantage. If it can create better three-point opportunities, the total expected points can still be higher. However, the Houston Rockets' 27 consecutive missed three-pointers showed me that expected value does not mean that the expected result will happen in every game. There is still variability, and an unlikely result can happen in one important game.

Another thing I learned is that expected value cannot explain everything in basketball. Statistics such as eFG%, TS%, turnovers, offensive rebounds, and free throws also affect offensive efficiency. This is why simply taking more three-pointers is not always the best strategy.

Overall, my calculations and the NBA examples helped me understand why three-point shooting has become so common. Three-point shots can be mathematically advantageous even with a lower shooting percentage because each successful shot is worth more points. At the same time, teams have to consider shooting percentage, shot selection, and the risk of missing. For me, this was a good example of how a relatively simple idea from probability can be used to understand something that happens in a real basketball game.

Works Cited

Oliver, D. (2004). Basketball on Paper: Rules and Tools for Performance Analysis. Potomac Books.

Kubatko, J., Oliver, D., Pelton, K., & Rosenbaum, D. T. (2007). “A Starting Point for Analyzing Basketball Statistics.” Journal of Quantitative Analysis in Sports, 3(3).

NBA. NBA Advanced Statistics. — eFG%, TS%, team and player shooting statistics.

NBA. Boston Celtics vs. Los Angeles Clippers, November 25, 2024, game statistics.

NBA. Golden State Warriors vs. Phoenix Suns, February 6, 2026, game statistics.

NBA. Houston Rockets vs. Golden State Warriors, Game 7, 2018 Western Conference Finals, May 28, 2018.

Goldsberry, K. (2019). SprawlBall: A Visual Tour of the New Era of the NBA. Houghton Mifflin Harcourt.

Sustaining the Ski Slopes: Snow Preservation through Biodegradable Geotextiles
By Connor Mason

Abstract

Rising global temperatures threaten snow cover at ski resorts worldwide, and the synthetic geotextile coverings that resorts use to slow melting shed microplastics into sensitive alpine environments. Biodegradable geotextiles offer a potential alternative, but their performance as snow coverings has not been compared beyond a single poly-lactic acid variant. This study compared four plant-fiber biodegradable geotextiles (burlap jute, coconut coir, hemp, and linen) against two controls (no covering and a synthetic nonwoven geotextile) in a controlled laboratory setting. Snow samples of standardized mass, density, and geometry (229 g; 350 kg/m^3) were placed on a 2:5-scale snow melt plate, covered, and allowed to melt at room temperature under indirect artificial light for 24 hours. Runoff mass was recorded every 5 minutes with a force sensor ($n = 3$ per group), normalized, and summarized as the area under the melt curve (AUC). A one-way ANOVA indicated a significant effect of covering type on AUC ($F(5, 12) = 6.616$, $p = 0.003$), and Tukey HSD post-hoc tests showed that hemp (mean AUC = 17.526) slowed melting significantly more than jute, coconut coir, the synthetic control, and the uncovered control; no other pairwise differences were significant. The hypothesis that linen would perform best because of its high albedo was not supported, most likely because insulation rather than reflectivity governed melt under indoor conditions. These results support including hemp in future field trials of biodegradable snow coverings as a microplastic-free alternative to synthetic geotextiles.

1. Introduction

On February 3rd, 2022, the first skiers took to the slopes for the Beijing Winter Olympics, marking the start of what would become a two-week extravaganza of winter sports (Pratt and Binner). But this time, something was different. Instead of the powdery, crystalline snow that most athletes were familiar with, all of the snow was fake (Guzman). For the first time in Olympic history, the Beijing Olympics relied entirely on artificially-manufactured snow, stepping away from the longstanding tradition of letting nature, at least partially, provide snow for the games (Guzman).

Although this may seem unexpected for those unfamiliar with the sport, to ski racers and enthusiasts, it's no surprise. Ever since the pre-industrial period, the Earth's global temperature has risen more than 1.5℃, bringing the planet to its warmest point in more than 11,000 years (Hansen et al.). From New South Wales, Australia, to the Swiss Alps, rising temperatures across the globe have had dramatic effects on the ski and snowboarding industry, forcing many resorts to seek alternative solutions to continue to sustain the snow (Pickering and Buckley).

As climate change transforms the global temperatures of the Earth into territories unheard of in decades prior, the once-stable ski industry faces risks that are difficult to assess (Steiger et al.). Based on a simulated model by researchers Mitterwallner, Steinbauer, Mathes, and Walentowitz, if climate change continues to progress at its current rate, within the century 13%

of global ski resorts will lose all snow cover, 20% will lose between 50% and 100%, and countless others will see similarly devastating impacts (Mitterwallner et al.). Even more alarming, "climate change has already decreased snowpack depths and reduced the amount of winter precipitation falling as snow," emphasizing that the climate change crisis isn't just a potential future concern for ski resorts, it's already here (Wilkins et al.).

2. Literature Review

2.1. Key Factors in the Snow Cover Problem

While the loss of snow cover may seem to be a problem only confined to the slopes themselves, the consequences of climate change have intense impacts on the environment and resort owners, both of which must be considered in evaluating potential solutions. For one, the existence of ski resorts is heavily intertwined with the environment and nature, making it vital that further damage is not done in the process of preserving ski resorts. The mountain regions where ski resorts are located are uniquely isolated from traditional human settlement, with ecology researchers Rixen and Rolando even considering them to be the "last strongholds of natural and seminatural terrestrial ecosystems." These ecological strongholds are also widely known to harbor a variety of species and rich biodiversity due to their unique combination of elevation, terrain, and environmental conditions found in few other places on the planet (Martin, qtd. in Rixen and Rolando). While essential, many present-day solutions seem to overlook the importance of ecological preservation.

As environmentalists push for options that don't harm the local environment, they often come into direct conflict with resort owners and operators, who prioritize long-term economic viability. Part of this is a consequence of the capital-heavy nature of the industry, since chair lifts, snow groomers, and terrain maintenance are all large, fixed-cost investments (Moreno-Gené et al., "Economic Sustainability"). At one ski resort in Austria, the constant cost of capital became so constricting that not even "1 cent of dividend" could be paid to investors across a multi-year period (Moreno-Gené et al., "Snow Tourism"). This budget compression highlights the necessity that solutions to the climate-induced snow issue are cost-effective, since many resorts simply cannot afford to purchase the complex systems that many current solutions demand.

2.2. Attempted Solutions to Maintain Snow Cover

Given the precarious economic standing of ski resorts today, many resort owners have attempted implementing solutions to keep snow, and paying customers, on the mountains. Out of the many ideas proposed, most resorts have opted for one of three solutions: artificial snowmaking, snow pits, or synthetic geotextiles (Weiss et al.).

For one, many resorts have chosen artificial snowmaking, a process where large snow fans combine high-pressure air and water to make snow (Zhao et al.). This allows resorts to counteract the loss of snow that climate change causes, an essential measure at planned competition events like the Winter Olympics (Zhao et al.). However, snowmaking has become more difficult in the past few years because it generally requires temperatures below 0℃,

something that “is becoming rarer” as temperatures rise (Cognard et al., “Snow Must Go On”). The solution is further held back by its resource-heavy nature, as the immense amounts of water required for snowmaking strain local water reservoirs, and consequently “[dam] rivers and [drain] wetlands” that the reservoirs pull water from (Cognard et al., “Snowmaking’s Slippery Slope”). The downsides are compounded by the lack of an economic upside to snowmaking, as Cognard et al. found even resorts who invested in snowmaking saw “no significant effect” on their profits (“Snow Must Go On”). With environmental implications and no long-term economic benefits, snowmaking is difficult to justify as a permanent solution.

Others have turned to snow pits, a solution which leverages the abilities of snow fans to create a base layer, or “pit,” of snow for early-season operations (Schmitt et al.). Through this solution, resorts artificially produce snow “under favorable ambient conditions in the winter” before piling the snow, covering it with sawdust or woodchips, and storing it outdoors over the summer months (Schmitt et al.; Grünewald et al.). Resorts can then redistribute the snow accumulated in snow pits across their skiable terrain early on in the season, providing the base coverage needed to begin operations sooner. While owners may appreciate this added revenue, this solution still falls short on environmental grounds. By creating such heavily compacted areas of snow, research shows that snow pits result in higher levels of plant mortality, decreased soil biological activity, and limited species recolonization even after the snow was removed (Buttler et al.).

In considering the known downsides of both artificial snowmaking and snow pits, some resorts have turned towards synthetic geotextiles, which are purpose-designed snow coverings that “reflect sunlight and reduce heat absorption, thereby slowing down the melting process” (Grau). Essentially, their light colors mean that they reflect energy from the sun, while their insulative properties serve as a shield against heat energy from the surrounding air. As energy is what causes snow to melt, reflecting and shielding against such energy absorption allows geotextiles to slow down the melting process (Grau; Zumdahl and DeCoste). At some resorts, synthetic geotextiles are used in tandem with snow piles, with sawdust and woodchips swapped out for geotextile coverings (Grau). At high-elevation resorts with year-round operations, this keeps snow on the mountains during the months where inconsistent temperatures may diminish the total snow cover (Senese et al.).

In many past studies, the implementation of geotextiles has been highly successful, with some finding that their use can slow the snow melting process by as much as 60% (Xie et al.). At the same time, geotextiles struggle with their environmental impact, as many common synthetic materials for geotextiles are known to degrade over time and result in “the accumulation of microplastics in the surrounding environment” (Elsayed et al.). This is concerning as ski resorts are often placed in some of the most pristine mountainous environments, which are home to “many plants and endemic species” (Moscovici).

2.3. Research Gap and Purpose

While geotextiles seem to be a promising option that balances the economic benefit of earlier opening dates with real-world practicality, these environmental concerns have hampered

their widespread adoption. In turn, some resorts have opted for biodegradable geotextiles, which package the benefits of the geotextiles in an environmentally-friendly, non-toxic format (Tanasă et al.). These are particularly useful because they break down, or biodegrade, when exposed to microorganisms present in soil, so any leftover fabric covering will decompose naturally over time (Brunšek et al.). Though synthetically manufactured biodegradable geotextiles exist, plant fiber-based variations are most commonly used today due to their abundance and reduced costs (Tanasă et al.).

While biodegradable geotextiles aren't a new development, previous research has centered around their uses in soil erosion, road construction, and drainage systems, as they benefit directly from the fabrics' degradable nature (Tanasă et al.). Existing research has yet to meaningfully extend the use of these fabrics to snow preservation and the ski industry specifically. Besides, the main study testing a biodegradable geotextile's impact on snow melt only considered one variant, poly-lactic acid (Senese et al.). With the exception of Senese et al., present studies which assess geotextile performance from a snow protection lens focus solely on synthetic, non-biodegradable covering materials like polypropylene (Grau; Xie et al.; Fischer et al.). With such a limited scope of research, it is generally unclear as to how different biodegradable geotextiles compare in performance as snow coverings. Consequently, scientists and resort owners have no indication of whether or not these environmentally-friendly geotextiles make effective coverings, and which may perform the most favorably if implemented. This study plans to fill this gap by comparing different biodegradable geotextiles and their performances as snow coverings.

This experiment investigates the research question, "How do biodegradable geotextile materials compare in their ability to reduce snow melt in a controlled laboratory environment?"Among the selected biodegradable geotextile options, materials vary significantly in natural color, with linen often being pale white, hemp being yellow, jute being golden-brown, and coconut coir being brown (EnviroTextiles). Based on these colors, it is hypothesized that if piles of snow are covered by a linen biodegradable geotextile, the snow will melt at a slower rate overall compared to uncovered piles of snow and other biodegradable materials. In theory, linen's white color, and hence, high light reflectivity (or albedo), will help it limit energy absorption from solar radiation, making it a more effective snow covering.

3. Methodology

3.1. Experimental Design

This study aims to compare the ability to which biodegradable geotextiles can abate snow melting through the analysis of burlap jute, coconut coir, hemp, linen, and two controls, no covering and a synthetic geotextile. These fabrics were selected because they had been specifically designated as biodegradable geotextiles by past researchers, and their unique densities, porosities, colors, and albedos[1] allow for multiple properties to be evaluated with each

[1] Albedo is a measure of how much light reflects off the surface of an object. Surfaces with higher albedos, or more light reflectivity, absorb less energy from the light source.

material (Tanasă et al.). Moreover, their popularity as fabrics means that each can already be purchased in bulk or wholesale quantities for large-scale applications. A quantitative laboratory experiment was set up which assessed how a snow sample covered by a geotextile melted over a period of 24 hours. As a whole, the procedure and setup for this study was heavily based upon Lintzén et al.'s work to study synthetic geotextiles as snow coverings in a laboratory. While field studies are often conducted for research testing geotextile snow coverings, this requires a high-elevation snow field which was unavailable for this study (Senese et al.; Nestler et al.). Additionally, field studies often struggle to control for changes in wind, sunlight, and temperature, making a laboratory study ideal (Lintzén et al.). Each experimental trial, which tested a single geotextile material, attempted to identify how the involved snow sample melted into water as a function of time. Experimental groups were dictated by fabric type, with three trials run per group (n=3) run to maximize the results' statistical significance. While n=3, or 18 total trials, may seem low, it was justified through Lintzén et al.'s study, who used a single trial per group.

3.2. Materials and Preparation

To begin, the snow that would be used in the study had to be collected. After snowfall the night prior, snow was collected on January 14th and January 29th, ensuring that the snow utilized was new, which gave it a consistent structure and density. In line with past studies that utilized homogenous layers of snow, all snow was taken from the same location, in areas free from footprints, animal tracks, or discoloration (Sergent et al.). In total, 18 bags of natural snow were collected, providing more than enough to be used for repeated trials of each fabric. The bags were stored in a freezer kept at -2℃.

Four biodegradable geotextiles, identified as commonly used by Tanasă et al., along with one synthetic geotextile control were purchased with the specifications shown in Figure 1. Each fabric chosen was 100% natural to ensure complete biodegradability. Six identical circular snow coverings (r=140mm) were cut from each material, giving every covering an equal area. Since each covering was congruent, material type was isolated as the independent variable. The thickness of each fabric was measured using a caliper while weights were provided by the manufactures; see Figure 1 for exact material specifications.

Fabric Type	Manufacturer	Specifications	Thickness	Cost per Sq Yard
Burlap Jute	EBP Inc.	72″ by 36″, 10oz/yd²	1.0mm	$6.50
Coconut Coir	Yarrdfir	12″ by 40″ (2x), 16oz/yd²	3.0mm	$35.11
Linen	Pllieay	39″ by 55″, 6.5oz/yd²	0.5mm	$12.08
Hemp	Hemp Go Green	36″ by 59″, 8oz/yd²	0.5mm	$12.20
Plastic Synthetic Geotextile	VEVOR	72″ by 1200″, 8oz/yd²	0.5mm	$1.09

Figure 1: Specifications for each of the fabrics purchased, including the synthetic control.

A snow melt plate was constructed, modeled after Lintzén et al., as shown in Figure 2. For the purposes of this study, a 2:5 scale was used, with a Styrofoam plate (shown in blue) and wood supports (shown in black). The 9° angle was kept the same since angles don't change with size, and triangular prism guardrails were made from plastic. The snow melt plate can be seen in its completed form in Figure 3.

To ensure that each snow sample had a consistent shape and geometry, a snow mold was needed. For the purposes of this study, a spherical cap-shaped mold, at a 2:5 scale of Lintzén et al.'s, was made out of poly-lactic acid (PLA) via 3D printing (shown in Figure 3). The same mold was used in each trial to maintain geometric consistency, which had a base diameter of 80mm and height of 56mm. This gave the spherical cap a volume of 655,926mm³ or 0.655L.

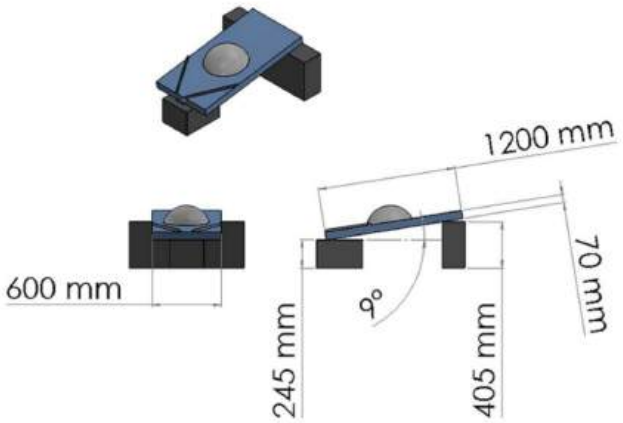


Figure 2: A digital model of the snow melt plate at full scale, reprinted from Lintzén et al.

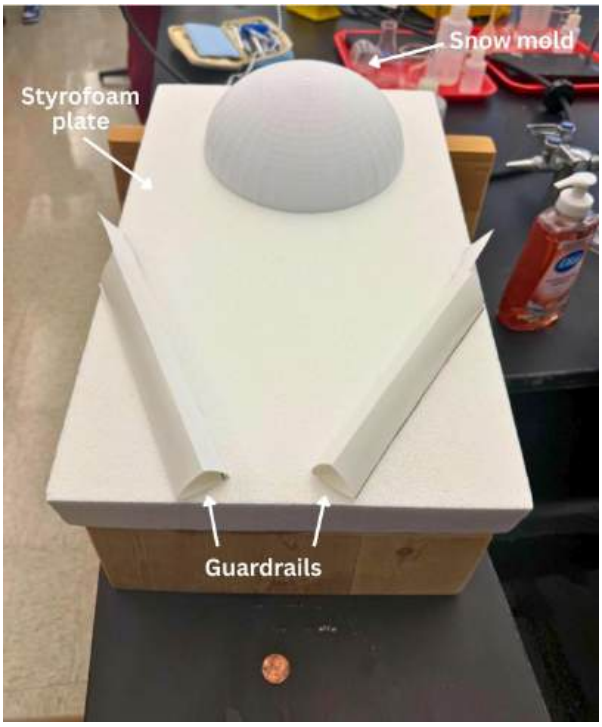


Figure 3: The 2:5 scale snow melt plate used in this study with labels.

3.3. Experimental Procedure

Prior to beginning trials, an experimental order was created. Since the conditions of the snow likely changed throughout the two-month testing window, trials run at the start would potentially be using slightly different snow than the ones at the very end. As a result, trial order was fully randomized to mitigate bias via the degradation of snow in storage.

At random, a bag of snow was chosen and removed from the freezer. To break apart any large clusters, the bag selected was pounded with a stainless steel scoop for 15 seconds. The snow mold (shown flipped over in Figure 3) was zeroed on a scale, and 229g of snow were scooped out of the bag and placed in the mold. Since the snow mold had a volume of 0.655L, using 229g of snow gave each sample a density of 350 kg/m^3, previously identified as a high average for natural snow (Judson and Doesken). Snow was brought to the snow melt plate, the mold was flipped onto the surface of the Styrofoam plate, and gently tapped three times to loosen the snow. The mold was carefully lifted up and any snow left in the mold was scraped onto the melt plate. Thanks to the mold, each trial began with a snow sample of the same approximate dimensions and shape. Once transfer was complete, the sample was covered with the designated fabric and data collection was begun. The human error in transferring snow was factored into the data through normalization; see 3.5: Data Analysis.

From a thermodynamics[2] standpoint, snow melts when it takes in enough energy to break apart the intermolecular forces[3] that bring molecules of H_2O together as a solid (Zumdahl and DeCoste). While this energy can come from a variety of sources, air temperature and radiation (generally from a light source) are seen as two of the most common for ski resorts (Xie et al.; Weiss et al.). To minimize bias from these factors, each trial was conducted in the same room at room temperature (21.4℃ - 24.1℃), while only an indirect, artificial light source was used (lightbulb: 32 Watts, 4100 Kelvin). Taking such precautions helped isolate fabric type as the primary independent variable.

3.4. Data Collection

Snow melt was determined by collecting the water from the melted snow in a pail, a similar system to Lintzén et al. The guardrails on the snow melt plate funneled the runoff down to a pail, which was attached to a force sensor hanging via a ring stand (see Figure 4). Using Vernier Graphical Analysis, the force sensor was set up to collect data every 5 minutes for 24 hours, generating 289 unique data points per trial. 5 minutes was chosen to compile data sets that reflected the full scope of the snow melting process and its intricacies, while 24 hours represented half of Lintzén et al.'s 48 hour period (whose setup was replicated at a 2:5 scale). Continuous data collection was used to understand the full melting process while avoiding manual collection for thousands of data points across all trials.



Figure 4: The complete snow melt plate and data collection setup.

[2] Thermodynamics is a branch of chemistry which focuses on energy, heat, and temperature. (Zumdahl and DeCoste).

[3] Intermolecular forces hold different molecules, like H_2O, close together as a solid. They differ from intramolecular forces, which bond together the hydrogen and oxygen within H_2O (Zumdahl and DeCoste),

3.5. Data Analysis

The Vernier force sensor, as mentioned above, collected force over time data across a 24 hour period in intervals of 5 minutes. The forces measured were converted to water runoff mass measurements through Newton's Second Law, shown in Figure 4 below.

$$F = m \times g$$

F: force of gravity (N), via force sensor

m: mass of water runoff (kg)

g: gravitational acceleration (9.8m/s²)

Due to human error in transferring the snow from the mold to the melt plate, each trial tested a slightly different amount of snow. To allow for direct comparison across trials, data from each trial was normalized from 0 to 1, with the max runoff mass during a trial equated to 1, a mass half of the max equated to 0.5, and so on for all of the data points. This approach mirrors common practice for normalization when direct comparison of raw data is not feasible as a result of collection methods (Tkachenko and Niedzielski).

After, the area under each melting curve (AUC) was estimated through the Trapezoidal rule[4]. Since the melting curves were defined by time (hours, 0-24) on the x-axis and normalized mass (0-1) on the y-axis, a trial where all of the snow melted instantaneously would produce an AUC of 24, while a trial where none of the snow melted would give an AUC of 0. By extension, trials which had an effective snow covering would produce a lower AUC, while trials with a less effective covering would have an AUC closer to 24. Evaporation of the runoff water, while present, was assumed to be inconsequential since it occurred consistently across all trials and the force data was normalized relative to the highest force.

Area under the curve was chosen as the primary metric for assessing the trials rather than a single point (like final runoff or melt time) to ensure that the results reflected the full melting process, which is often nonlinear. Using AUC avoids misclassifying materials as performing favorably if, for instance, they cause 90% of the snow to melt very quickly and abate the remaining 10% from melting for several hours.

4. Results

To compare each material, descriptive statistics were created for each group tested (biodegradable: hemp, jute, linen, coconut coir; control: no fabric, synthetic geotextile) and their respective AUCs. The mean and standard deviation AUC values can be shown in Figure 5. Based on the results in Figure 5 and Figure 6, hemp produced the lowest mean AUC value, implying that it had the most success in extending the snow melting process. By contrast, no fabric (control) had the highest mean AUC, indicating that all coverings slowed the melting process compared to what's normal.

[4] Trapezoidal rule is a Calculus I technique which estimates the area under a curve by making many small trapezoids and summing their areas. This was needed since the exact area under a curve cannot be calculated through data points alone.

Total Normalized Area Under the Curve (water runoff (kg) / time)						
	Burlap Jute	Coconut Coir	Hemp	Linen	No Fabric (Control)	Nonwoven Geotextile (Control)
Mean AUC	20.167	19.487	17.526	19.251	20.340	20.135
St. Deviation AUC	0.167	0.835	1.169	0.208	0.089	0.737

Figure 5: The calculated mean, median, and standard deviation AUC for each fabric type.

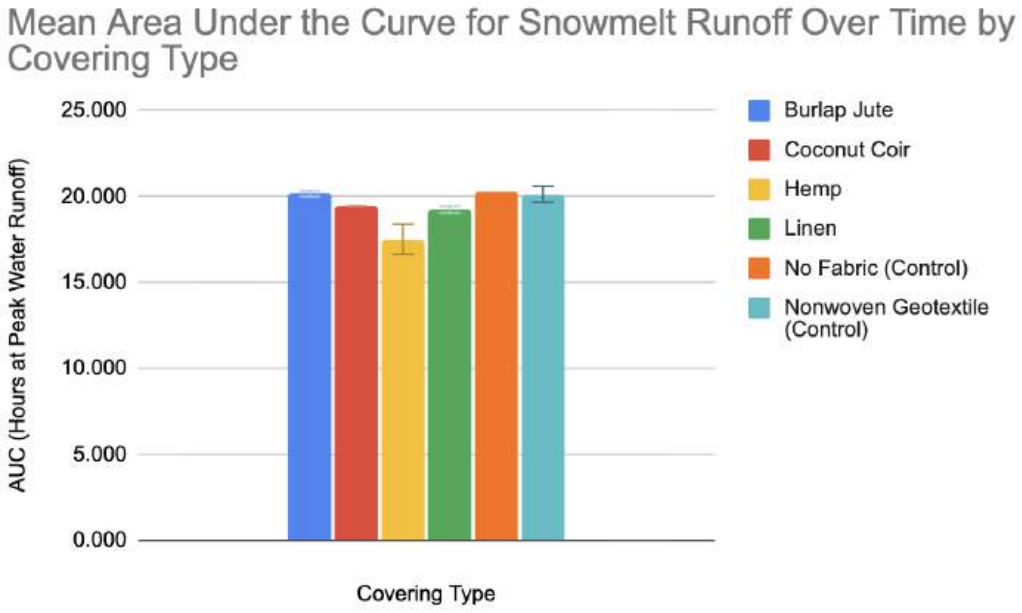


Figure 6: A graph displaying the mean and standard deviation AUC for each fabric type.

To analyze the results of this study, an ANOVA test was conducted to determine if the different covering types differed significantly in their total AUC for the normalized force over time graph. Based on the analysis, it can be concluded that there is a significant difference between covering type and AUC (FStat |6.616| > FCrit |3.105|, P-value 0.003 < alpha 0.05). However, the ANOVA test only concluded that there is a significant difference between some of the groups, but does not indicate which groups, specifically, differ in their AUC. To understand the groups which differ significantly from one another, a Tukey HSD post-hoc test was performed. The Tukey HSD post-hoc analysis revealed that hemp produced a statistically significant lower AUC when compared to burlap jute (pCalc 0.006 < pCrit 0.05), coconut coir (pCalc 0.047 < pCrit 0.05), the nonwoven geotextile control (pCalc 0.007 < pCrit 0.05), and the no fabric control (pCalc .004 < pCrit 0.05). None of the other groups produced significant differences in results. At large, hemp was the only material that differed significantly in AUC from any other group, while the rest of the groups, when paired together, were not statistically significant.

5. Discussion

5.1. Interpretation of Findings

This study sought to compare the ability of biodegradable geotextile materials to reduce snow melt in a laboratory environment. Through a laboratory study centered around an indirect light source and above-freezing ambient temperatures, preliminary results can be identified to understand the way that the properties of biodegradable and synthetic covering materials influenced snow melt. Given that all trials were conducted indoors, solar and UV radiation were minimized. This left air temperature as the main source of energy for snow melt in this experiment. Geotextiles primarily combat this ambient temperature by insulating the snow and creating a physical barrier between the snow and the air, so results should be considered in the context of insulation (Grau).

The suggestion that a hemp covering most effectively slowed the process of snow melt is intriguing, as it conflicts with a variety of the material properties generally correlated with insulation. Hemp had very similar properties to the synthetic plastic geotextile (control), with both sharing a weight of 8oz/yd² and thickness of 0.5mm per Figure 1. Based on the similar thickness, one may assume that they would insulate and shield against ambient heat-driven melting in a similar manner, yet hemp's mean AUC was lower than the synthetic geotextile, at 17.526 and 20.135, respectively (see Figure 5).

Moreover, compared to the rest of the biodegradable geotextiles, hemp has relatively similar material properties; hemp's thickness of 0.5mm is consistent with linen and lower than burlap jute and coconut coir while its weight of 8oz/yd² places it firmly between the lowest of 6.5oz/yd² for linen and the highest of 16oz/yd² for coconut coir (see Figure 1). Nevertheless, hemp was the only biodegradable geotextile that was statistically different from the controls, suggesting that factors outside of thickness and weight may have driven insulation and effectiveness. One outstanding variable between groups is fiber composition[5], which varied due to the natural plant structure of each material. By theoretically creating different spaces for air pockets within the coverings, fiber composition could have impacted the insulative qualities of each material, for better or worse. Still, this remains a speculation and further research on fiber composition within geotextiles is needed.

Prior to beginning trials, it was hypothesized that linen would be the most effective biodegradable covering because of its light color and presumed high albedo. This hypothesis was grounded in previous field research on snow coverings, which found that the materials with the highest albedos were the most effective snow coverings since they protected against UV and solar radiation (Xie et al.; Senese et al.). However, research conducted after making this hypothesis revealed that insulation, not albedo, is the primary driving factor for geotextiles' effectiveness indoors (Lintzén et al.). With only marginal amounts of UV radiation from the lightbulb in the room, insulation, not albedo, is likely what helped the coverings slow the melt

[5] Fiber composition refers to the unique arrangement of fibers, which are thread-like filaments, that can give materials structure, strength, and flexibility (Hu et al.).

process (Xie et al.). Correspondingly, the results of this study didn't support the hypothesis, since the hypothesis centered around a means of melting which was ultimately not present in this study.

Although a past study identified biodegradable geotextiles as less effective snow coverings than synthetic ones, this study's results found their performance to be similar or better than a synthetic control (Senese et al.). One explanation could be Senese et al.'s use of the biodegradable plastic poly-lactic acid, which lacks the natural fiber composition present in fabrics like hemp and linen. The lack of these fibers could have hindered the material's success, though more research is needed to know for sure.

5.2. Implications

Based on the data and results collected from this laboratory study, biodegradable geotextiles, and hemp particularly, can work as snow coverings to slow snow melting. This justifies hemp's place in future field studies comparing geotextile options, as it was the strongest performer in a laboratory-based experiment with 4 different natural fiber coverings. These field studies could provide more definitive results regarding the viability of biodegradable geotextiles as an environmentally-friendly, microplastic-free alternative to synthetics, while properly analyzing how both albedo and insulation impact effectiveness.

Practically, resorts may consider beginning their own small-scale field tests for different covering types, helping to understand what works the best on an individual basis. As factors like wind, sunlight, shade, and altitude may influence a geotextile's effectiveness in a unique way, having these small-scale studies could help identify the best covering for a resort's local attributes.

Ski resort owners will also need to consider the cost to implement these potential solutions. This study's acquisition cost for the fabrics can be seen in Figure 1. Per the table, a 12 fold difference can be observed between hemp, the strongest biodegradable geotextile, and the synthetic geotextile control. While the actual acquisition cost may be lower for resorts looking to purchase fabric at a large scale, the implementation of biodegradable geotextiles on environmental grounds may only make sense if it is similarly or more cost-effective than alternatives like synthetic geotextiles, snowmaking, and snow pits.

5.3. Limitations

To begin, the study was conducted in a laboratory environment, which meant that UV and solar radiation were not present as they would have been in a field study. By excluding these two potential energy sources, the albedo of each fabric, or its ability to reflect solar radiation, did not meaningfully factor into the results. In consequence, the trials disadvantaged materials with high albedos (or light colors) since their inherent advantages were controlled for. In a field study, those fabrics may be perceived as more effective because their albedos could help them reflect energy away from the snow.

Next, the spherical cap shape of the snow mold was unnatural, forcing geotextiles to drape down significantly to cover the entire snow sample. While modeled after Lintzén et al. at a

2:5 scale, the snow's height of 56mm was substantial for coverings that had a radius of only 140mm. Accordingly, lightweight materials like linen may have had an easier time blocking out ambient thermal energy since they could drape over the entire snow sample, while heavier options like coconut coir were forced to adopt a u-shape that left some snow uncovered (see coconut coir in Figure 4 as an example of this). This may have favored lighter fabrics to perform better, since they could more fully cover the samples of snow.

Lastly, each fabric was only tested three times, which may have allowed external variables to impact the effectiveness, or lack thereof, of the coverings. Ironically, despite being the only material with statistically significant results, hemp had the highest standard deviation at 1.169, suggesting that even its results could have been swayed by external influences (see Figure 5). Had more tests been conducted for each experimental group, the impact of external variables would have been reduced, giving way for more reliable comparisons between the different types of geotextiles.

5.4. Future Research

To continue investigating the potential of biodegradable geotextiles as an alternative to synthetics, further research is needed. Since the results of this study are specific only to laboratory conditions, researchers should perform a comparative analysis of biodegradable geotextiles in a field, or outdoor, format. This would provide a more nuanced understanding of how surface reflectivity and albedo support the effectiveness of certain coverings, as a lack of solar and UV radiation meant that neither were considered effectively in this study. By conducting such a study at a large scale, researchers could also account for the spherical cap drape and sample size limitations of this study. Future researchers should also consider testing synthetic biodegradable geotextiles, like poly-lactic acid and Secutex® Green, in comparison to the plant-based biodegradable geotextiles used in this study, to understand which is a more realistic material for widespread adoption. One final avenue for future research is small-scale geotextile tests conducted on-site at ski resorts, including both synthetic and biodegradable variants, giving insight into how geographical factors like wind, elevation, and sunlight affect the effectiveness of the coverings. By mapping the results of these studies geographically, researchers could learn how regional differences in climate favor specific geotextile materials over others. While there's undeniably a long road ahead, this future research just might reveal the key to sustaining the ski slopes.

Acknowledgments

I give my thanks to Mark Wollschlaeger, AP Research Teacher at Hinsdale Central High School, for his guidance throughout the design and execution of this study.

Works Cited

Brunšek, R., et al. "Biodegradation Properties of Cellulose Fibers and PLA Biopolymer." Polymers, vol. 15, no. 17, 2023, p. 3532, https://doi.org/10.3390/polym15173532.

Buttler, A., et al. "Impacts of Snow-Farming on Alpine Soil and Vegetation: A Case Study from the Swiss Alps." Science of the Total Environment, vol. 903, no. 1, 2023, p. 166225, https://doi.org/10.1016/j.scitotenv.2023.166225.

Cognard, J., et al. "Snowmaking's Slippery Slope: The Effect of Mountain Reservoirs on Water Demand." Ecological Economics, vol. 233, no. 1, 2025, p. 108586, https://doi.org/10.1016/j.ecolecon.2025.108586.

Cognard, J., et al. "The Snow Must Go On: Can Snowmaking Keep Ski Resorts Profitable in a Changing Climate?" Current Issues in Tourism, vol. 28, no. 24, 2024, pp. 1–18, https://doi.org/10.1080/13683500.2024.2409862.

Elsayed, H., et al. "Geotextiles for Environmental Protection, a Review." Egyptian Journal of Chemistry, vol. 69, no. 2, 2026, https://doi.org/10.21608/ejchem.2025.402197.12025.

EnviroTextiles Editors. EnviroTextiles®: Sustainable Fabrics from Hemp & Other Green Textiles. Envirotextiles, 2024, https://www.envirotextiles.com/.

Fischer, A., et al. "Local Reduction of Decadal Glacier Thickness Loss through Mass Balance Management in Ski Resorts." The Cryosphere, vol. 10, no. 6, 2016.

Grau, O. Innovative Snow Storage Solutions for the Pyrenees Ski Resorts. 2024. Bachelor's thesis.

Grünewald, T., et al. "Snow Farming: Conserving Snow over the Summer Season." The Cryosphere, vol. 12, no. 1, 2018, pp. 385–400, https://doi.org/10.5194/tc-12-385-2018.

Guzman, C. "What Artificial Snow at the 2022 Olympics Means for the Future of Winter Games." Time, 8 Feb. 2022, https://time.com/6146039/artificial-snow-2022-olympics-beijing/.

Hansen, J. E., et al. "Global Warming Has Accelerated: Are the United Nations and the Public Well-Informed?" Environment: Science and Policy for Sustainable Development, vol. 67, no. 1, 2025, pp. 6–44, https://doi.org/10.1080/00139157.2025.2434494.

Hu, J., et al. Handbook of Fibrous Materials. Wiley Online Library, 2020.

Judson, A., and N. Doesken. "Density of Freshly Fallen Snow in the Central Rocky Mountains." Bulletin of the American Meteorological Society, vol. 81, no. 7, 2000.

Lintzén, N., et al. "Laboratory Investigation of Different Insulating Materials Used for Snow Storage." Journal of Cold Regions Engineering, vol. 33, no. 4, 2019, https://doi.org/10.1061/(ASCE)CR.1943-5495.0000194.

Mitterwallner, V., et al. "Global Reduction of Snow Cover in Ski Areas under Climate Change." PLOS ONE, vol. 19, no. 3, 2024, e0299735, https://doi.org/10.1371/journal.pone.0299735.

Moreno-Gené, J., et al. "The Economic Sustainability of Snow Tourism: The Case of Ski Resorts in Austria, France, and Italy." Sustainability, vol. 10, no. 9, 2018, p. 3012, https://doi.org/10.3390/su10093012.

Moreno-Gené, J., et al. “Snow Tourism and Economic Sustainability: The Financial Situation of Ski Resorts in Spain.” Applied Economics, vol. 52, no. 52, 2020, pp. 1–19, https://doi.org/10.1080/00036846.2020.1770683.

Moscovici, D. “Ski Resort Closures and Opportunities for Sustainability in North America.” Land, vol. 25, no. 11, 2022, https://doi.org/10.3390/land11040494.

Nestler, A., et al. “Hydrological Implications of Covering Wind-Blown Snow Accumulations with Geotextiles on Mount Aragats, Armenia.” Geosciences, vol. 4, no. 3, 2014, pp. 73–92, https://doi.org/10.3390/geosciences4030073.

Pickering, C. M., and R. C. Buckley. “Climate Response by the Ski Industry: The Shortcomings of Snowmaking for Australian Resorts.” AMBIO, vol. 39, no. 5–6, 2010, pp. 430–438, https://doi.org/10.1007/s13280-010-0039-y.

Pratt, J., and A. Binner. “Mikael Kingsbury Shows Class in Reaching Beijing 2022 Moguls Final.” Olympics.com, International Olympic Committee, 3 Feb. 2022, https://www.olympics.com/en/news/mikael-kingsbury-shows-class-in-reaching-beijing-2022-moguls-final.

Rixen, C., and A. Rolando. The Impacts of Skiing and Related Winter Recreational Activities on Mountain Environments. Bentham Science Publishers, 2013.

Schmitt, L., et al. “Life Cycle Assessment of Snow Farming: A Case Study on Climate Change Adaptation of Cross-Country Skiing Centers in Germany.” Journal of Sustainable Tourism, 2025, pp. 1–32, https://doi.org/10.1080/09669582.2025.2571957.

Senese, A., et al. “The Non-Woven Geotextiles as Strategies for Mitigating the Impacts of Climate Change on Glaciers.” Cold Regions Science and Technology, vol. 173, no. 1, 2020, p. 103007, https://doi.org/10.1016/j.coldregions.2020.103007.

Sergent, C., et al. “Experimental Investigation of Optical Snow Properties.” Annals of Glaciology, vol. 17, no. 1, 1993, https://doi.org/10.3189/S0260305500012970.

Steiger, R., et al. “A Critical Review of Climate Change Risk for Ski Tourism.” Current Issues in Tourism, vol. 22, no. 11, 2017, pp. 1343–1379, https://doi.org/10.1080/13683500.2017.1410110.

Tanasă, F., et al. “Geotextiles—A Versatile Tool for Environmental Sensitive Applications in Geotechnical Engineering.” Textiles, vol. 2, no. 2, 2022, pp. 189–208, https://doi.org/10.3390/textiles2020011.

Tkachenko, Y., and P. Niedzielski. “FTIR as a Method for Qualitative Assessment of Solid Samples in Geochemical Research: A Review.” Molecules, vol. 27, no. 24, 2022, p. 8846, https://doi.org/10.3390/molecules27248846.

Weiss, H. S., et al. “Optimization of Over-Summer Snow Storage at Midlatitudes and Low Elevation.” The Cryosphere, vol. 13, no. 12, 2019, pp. 3367–3382, https://doi.org/10.5194/tc-13-3367-2019.

Wilkins, E. J., et al. “Climate Change and Utah Ski Resorts: Impacts, Perceptions, and Adaptation Strategies.” Mountain Research and Development, vol. 41, no. 3, 2021, https://doi.org/10.1659/mrd-journal-d-20-00065.1.

Xie, Y.-D., et al. “Assessing the Impact of Artificial Geotextile Covers on Glacier Mass Balance and Energy Fluxes.” Advances in Climate Change Research, vol. 15, no. 6, 2024, https://doi.org/10.1016/j.accre.2024.10.007.

Zhao, W., et al. “An Experimental Investigation on the Size Distribution of Snow Particles during Artificial Snow Making.” Energies, vol. 16, no. 21, 2023, p. 7276, https://doi.org/10.3390/en16217276.

Zumdahl, S., and D. DeCoste. *Introductory Chemistry*. 9th ed., Cengage Learning, 2019.

Testing the 70-30 Rule: A Simulation-Based Evaluation of Aggressive and Conservative Play in Junior Girls' Golf By Kexin Zhang

Abstract

A new strategic approach known as the 70-30 Rule has recently gained attention in golf, suggesting that over a round, golfers should play approximately 70% of shots conservatively and the remaining 30% aggressively. The idea is that careful play reduces the chance of costly mistakes, while a few aggressive shots allow golfers to take advantage of high-reward scoring opportunities. This paper evaluates whether this strategy actually helps junior golfers. Thirty shots per club were collected from each of the three junior female golfers using the TrackMan golf simulator. We used their shot distributions to build simulations in Google Earth and the Desmos Graphing Calculator. We then scored the results using the Expected Points System we developed for this study. Finally, we analyzed the results with a two-way analysis of variance with replication. We found a significant effect of strategy on Expected Points ($p = 2.6306 \times 10^{-10}$). We also found a significant player × strategy interaction, but no single strategy was best for all three players. The 70-30 plan gave the lowest Expected Points for one player, the 60-40 plan for another, and the 80-20 plan for the third. These results do not support the 70-30 Rule as something that works for every golfer. They do suggest that the percentage of aggressive shots affects performance, and that the best balance may depend on the player. Our sample is small, so this study is best seen as a first look that future work can build on.

Introduction

When people think of sports, they often think about power, technique, and physical strength; however, a major yet often overlooked aspect is strategic planning. A 2024 review of heuristics in sport found that they matter in competition because they are quick and they work (Ramanayaka et al.). Heuristics are mental shortcuts. The review examined 38 studies and found that athletes often rely on these shortcuts when making decisions. Shortcuts help athletes make decisions quickly; numbers, however, can show which strategies are most likely to work. This is especially true in golf, where each shot presents several options with different levels of risk and reward. This review also noted that research on heuristics in sports remains limited, particularly regarding studies focused on women athletes.

Several published studies have investigated decision-making in golf through a qualitative approach. Interviews with male professional golfers suggest that golfers do not decide based only on statistics; experience and confidence matter too (Gnagy et al.). The researchers found that expert golfers make decisions based on experience and competitive context, not only on what might be the most "logical" choice. Amy Whitehead showed that the Think Aloud Method (TAM) is a valid way to study decision-making during play (Whitehead). In this method, golfers say their thoughts out loud while they play. Study 1 of this thesis compares the cognitive processes of 30 skilled and 30 novice golfers during a putting task. Whitehead and Patricia Jackman later investigated adolescent golfers' cognitive processes during competitive rounds

using the TAM (Whitehead and Jackman). They found six cognitive themes that repeatedly appeared, including how golfers judge risk and strategy.

Psychological factors also show up in performance data. Taylor Broughton compared golfers and non-golfers on three tests of risk-reward decision-making under stress (Broughton). The study concluded that golfers generally regulate stress more effectively and are willing to take more risks than non-golfers, although they did not outperform non-golfers on the tasks. This study highlights the role of cognitive and emotional factors in golfers' decision-making under pressure. Devin Pope and Maurice Schweitzer analyzed over 2.5 million putts from PGA Tour players and found that golfers do better on par putts than on birdie putts of the same difficulty (Pope and Schweitzer). They suggest that even elite players like Tiger Woods are affected by loss aversion, illustrating that psychological factors play a vital role in sports. Together, these studies help explain how and why golfers make certain decisions, but do not directly compare the effectiveness of course management strategies.

Furthermore, Sue Shapcott produced a USGA literature review that compiled and summarized existing research on how golfers make decisions during a round (Shapcott). This review served as a starting point for identifying relevant studies, which were examined alongside additional sources found independently.

Other research has shown that golf strategy can be modelled and tested; however, almost all of it uses data from elite male professionals. Michael Stöckl et al. collected data from the ShotLink database for all 2011 PGA Tour tournaments (Stöckl et al.). They developed the ISOPAR method to reassess hole difficulty beyond distance alone, enabling golfers to visualize the challenging areas of each hole relative to the field's average performance. Their finding that performance depends on more than distance motivated the present research to consider both shot dispersion and landing location when evaluating strategic outcomes. Nishad Wajge and Gautier Stauffer compared strategies in stroke play and match play (Wajge and Stauffer). They built mathematical and statistical models, used historical PGA Tour data to compare match-play strategies, and concluded that the advantage of tailoring strategy to one's opponent is, overall, very small. Additionally, Stauffer and Matthieu Guillot used a Markov Decision Process model to test the "Drive for show, putt for dough" adage (Stauffer and Guillot). They found that their results did not support the traditional emphasis on putting over driving. Both studies used quantitative approaches to evaluate golf strategy. Their studies demonstrate how data analysis can assess golf strategies. They found that the gain in strokes per hole was larger for driving than for putting, going against what many golfers believe. Our paper uses the same basic idea: take a golf adage and test it with a model.

Researchers also used distributions and computer programs to investigate decision-making in golf. Marie Hardt and Dan Nettleton modelled hole scores using the Hardy distribution, a probability distribution designed to model golf scores on individual holes. They compared PGA Tour professionals with female high-school golfers and found that scores varied more among the high-school players (Hardt and Nettleton). That group is close to the golfers in our study. Soonmin Ko uses a computer simulation to examine how varying levels of golf skill

affect scoring outcomes (Ko). The study draws an interesting connection between decision-making in golf and optimizing investment policies in finance. This paper further showcases the value of computational analysis in golf and supports simulation methods for estimating golfers' performance. Furthermore, Mehdi Khazaeli and Leili Javadpour developed an artificial intelligence system to help golfers choose the most suitable club for a shot (Khazaeli and Javadpour). They considered multiple factors, including distance, terrain, wind strength and direction, and lie quality, and the model recommends four possible club options for each shot. This study shows how artificial intelligence can be integrated into sports decision-making based on statistical models to deliver more effective recommendations for golfers. Recent work has also attempted to quantify course-management decisions.

David Percy developed a mathematical model to predict outcomes and guide decision-making in sports. This study applies stochastic process models (probability models that include randomness) to develop a mathematical framework for strategy selection when outcomes are uncertain (Percy). This model uses dynamic learning algorithms to update its predictions based on previous results. This study establishes a promising connection between mathematics and sports. Although not specific to golf, this study demonstrates that developing mathematical and statistical models to optimize decision-making is a common objective across many sports. Most relevant, Andrew Mack develops a theoretical model explaining how golfers should make strategic decisions to minimize their scores, applying economic and decision-theoretic concepts to golf while focusing on risk and expected outcomes (Mack). This study examines situations such as whether to play aggressively or conservatively, concluding that golfers should adopt strategies based on expected outcomes rather than purely on risk avoidance. That is the idea behind our Expected Points system.

A smaller group of studies has looked at risk-taking directly. Serkan Ozbeklik and Janet Smith investigated professional players' decision-making in match play (Ozbeklik and Smith). Specifically, they analyzed how golfers adjust their risk-taking behaviour when they are winning or losing. They found that decisions are closely tied to the match score, the number of holes remaining, and hole difficulty. Todd McFall and Kurt Rotthoff looked at stroke play instead. They analyzed data from 10 years of professional golf tournaments to examine how golfers adjust their risk-taking strategies (McFall and Rotthoff). In particular, they examined the second shot on par-5 holes, a key opportunity for risk-taking. They concluded that bigger rewards, tournament position (such as being near the cutline), or playing alongside "Superstar Tiger Woods" encourage professional golfers to take greater risks. Together, these studies demonstrate that risk-taking is an important component of golf strategy and that golfers change their decisions based on the rewards available and where they stand in a tournament. Our study builds on this. We examine whether golfers can plan the balance between aggressive and conservative play for an entire round.

The 70-30 Rule is similar to the old "Drive for show, putt for dough" adage, advice golfers pass around without testing. LuxGolf, a commercial golf website, published it at the end of 2025 (Admin). The 70-30 Rule states that, over the course of a round, golfers should play

approximately 70% of shots conservatively, while the remaining 30% should be played aggressively (Admin). The reasoning is that conservative play reduces the chance of detrimental mistakes, while selective aggressive shots allow golfers to capitalize on high-reward scoring opportunities. This approach provides golfers with a framework for developing a game plan—a pre-round strategy that outlines how to play each hole on a course. Adhering to this plan throughout a tournament round may help golfers minimize unnecessary risks while creating more favourable scoring opportunities. However, it remains unclear whether consistently following the 70-30 Rule improves performance or oversimplifies the game's strategic demands.

From the 2011 study by Pope and Schweitzer on golfers' tendency toward loss aversion to more recent articles evaluating strategies such as "Drive for show, putt for dough" and the 70-30 Rule, researchers have shown increasing interest in the psychological and strategic aspects of decision-making in golf. New strategic theories continue to be proposed to help golfers improve their performance. Because the 70-30 Rule was introduced recently, it has not been published in peer-reviewed research, so it is unclear whether following it truly improves performance or whether it oversimplifies strategy in golf. At the same time, junior sports participation has increased substantially in recent years. According to the National Golf Foundation, junior golf participation has risen by 58% in the United States since 2019 (National Golf Foundation). If this rule proves effective, it could help junior golfers further improve their performance by developing a game plan that maximizes their chances of success.

This paper tests the hypothesis that if 30% of shots in a round of golf are played aggressively, junior golfers will achieve a lower mean Expected Points total. Expected Points measures how good a golfer's ball positions are across a round. We collected data from three junior female golfers, a population for which research on strategic decision-making remains limited despite the continued growth of junior girls' golf; given the small participant size, the findings are intended as an initial evaluation rather than a widely generalizable conclusion. We used their shot distributions to build a simulation. We then analyzed the results with a two-way analysis of variance (ANOVA) with replication, using player and strategy as the two factors. Ultimately, this study seeks to determine whether the 70-30 Rule is an effective, evidence-based strategy for improving performance among junior female golfers similar to those examined in this study.

Methodology

This study employed a simulation-based experimental design to investigate how the percentage of shots played aggressively affects a round's overall Expected Points total. We compared play that was 20%, 30%, and 40% aggressive, using shot-performance data we collected from three junior golfers. These percentages provided a full-swing implementation of the 70-30 Rule, with approximately 30% of full-swing shots played aggressively under the 70-30 Rule. We used a two-way ANOVA with replication to assess whether Expected Points differed significantly across the three percentages.

This study included three Canadian junior female golfers aged 15 to 18. Their handicaps ranged from +1.5 to +3.1, and their heights from 168 cm to 179 cm. A handicap is a number computed from a golfer's recent scores that shows how well they play; lower handicaps, including plus handicaps, indicate a higher skill level. All three golfers competed at the international level, and their best AJGA rankings range from 80 to 400. We explained the study's purpose and procedures to all participants. They all participated voluntarily, with permission from a parent or guardian. During data collection, participants performed standard golf shots using their usual technique. We distributed the data collection across at least three separate sessions to reduce fatigue and minimize additional physical risk. We kept the participants anonymous throughout to protect privacy.

Each player hit 30 shots with every one of the 11 clubs in their bag, so we had the full range of their full-swing distances recorded. Players used the same clubs they normally compete with in tournaments, all of which comply with AJGA and Golf Canada equipment regulations. We collected the data indoors using TrackMan, a golf launch monitor and simulator. TrackMan uses Doppler radar and high-speed camera technology to track the club and ball through impact and flight. For indoor use, TrackMan measures variables such as club speed, ball speed, and launch angle, then uses these data to calculate the ball's projected flight and landing position. Professional golfers and the PGA Tour use TrackMan for shot tracking and performance analysis. We collected our data over three weeks to keep the sessions as consistent as possible. When finished, we exported the shot history from the TrackMan software and took out two measurements. Side is how far left or right the ball landed from the target line. Carry is how far the ball flew before it first hit the ground.



Figure 1: TrackMan interface displaying the shot analysis of a participant hitting a 9 iron. The software records launch parameters and calculates variables such as Carry and Side, which we exported for this statistical analysis.

We organized the data in tables, with Side in the left column and Carry in the right column. Side values were originally measured in feet, so we converted them to yards to keep all the units consistent. In practice, how far a ball rolls after it lands depends heavily on the course conditions. Wet or soft ground reduces roll, while dry or hard ground increases it. For every club except the driver, we used carry distance only, as these clubs typically experience relatively little roll after landing. The driver is different. It produces the fastest ball speed and the most roll, so carry distance alone does not realistically represent how far the shot would travel during play. For the driver, we therefore looked at two more measurements. Total is the distance from impact to where the ball finally stops, including roll. Total Side is how far left or right the ball ends up from the target line. We then averaged Carry with Total, and Side with Total Side, and used those two averages as the driver's distance and side. We made this adjustment because TrackMan usually predicts more roll than a golf course actually provides, and actual roll varies with fairway conditions. Averaging provides a more realistic estimate of the driver's final position.

			6.00				5.00				hybid			
2	ide (Yds)	Cay (Yds)	Ball #	Side (Ft)	Side (Yds)	Cay (Yds)	Ball #	Side (Ft)	Side (Yds)	Cay (Yds)	Ball #	Side (Ft)	Side (Yds)	Cay (Yds)
3	-1.40	138.40	1.00	13.40	4.47	152.70	1.00	-6.70	-2.23	159.20	1.00	-1.70	-0.57	163.60
4	8.70	136.00	2.00	26.00	8.67	147.40	2.00	9.90	3.30	158.10	2.00	-5.00	-1.67	162.50
5	-6.80	143.30	3.00	19.40	6.47	150.00	3.00	0.50	0.17	154.40	3.00	-27.10	-9.03	167.50
6	-8.40	143.90	4.00	-9.40	-3.13	145.40	4.00	23.60	7.87	159.40	4.00	-7.70	-2.57	163.50
7	-6.40	138.10	5.00	13.80	4.60	149.30	5.00	23.50	7.83	158.70	5.00	2.80	0.93	167.90
8	-1.30	138.70	6.00	-10.10	-3.37	151.40	6.00	10.20	3.40	161.60	6.00	24.60	8.20	166.30
9	5.05	140.20	7.00	-5.60	-1.87	152.10	7.00	-4.20	-1.40	157.10	7.00	-38.40	-12.80	167.40
10	2.80	137.00	8.00	-14.90	-4.97	153.00	8.00	-9.40	-3.13	153.10	8.00	-7.60	-2.53	169.30
11	-0.90	142.40	9.00	-18.10	-6.03	148.50	9.00	29.40	9.80	160.40	9.00	-12.90	-4.30	171.80
12	3.95	141.50	10.00	-5.80	-1.93	151.40	10.00	-22.70	-7.57	157.50	10.00	-8.50	-2.83	168.90
13	-0.50	141.80	11.00	-13.50	-4.50	152.50	11.00	9.40	3.13	159.50	11.00	-0.70	-0.23	172.20
14	-11.35	142.90	12.00	-5.60	-1.87	144.80	12.00	-10.40	-3.47	156.70	12.00	-1.40	-0.47	169.80
15	-0.10	141.80	13.00	-6.90	-2.30	153.20	13.00	6.40	2.13	161.10	13.00	11.90	3.97	173.40
16	-5.45	135.80	14.00	15.00	5.00	148.00	14.00	-11.60	-3.87	160.20	14.00	25.80	8.60	172.00
17	4.65	142.70	15.00	-1.30	-0.43	156.00	15.00	10.80	3.60	162.40	15.00	-20.70	-6.90	168.40
18	-10.70	143.50	16.00	-16.60	-5.53	151.50	16.00	0.20	0.07	160.20	16.00	12.30	4.10	175.50
19	-1.40	146.00	17.00	-24.70	-8.23	153.30	17.00	-1.90	-0.63	163.90	17.00	-21.70	-7.23	175.10
20	-0.20	143.20	18.00	16.00	5.33	151.70	18.00	0.30	0.10	160.80	18.00	-7.70	-2.57	177.00
21	-7.10	139.70	19.00	-16.50	-5.50	155.10	19.00	16.90	5.63	163.00	19.00	6.80	2.27	175.70
22	-1.95	143.20	20.00	6.20	2.07	150.40	20.00	11.80	3.93	160.90	20.00	4.10	1.37	161.30
23	-9.70	144.40	21.00	6.70	2.23	152.10	21.00	-3.80	-1.27	159.10	21.00	19.40	6.47	169.60
24	-7.45	141.60	22.00	14.90	4.97	155.30	22.00	12.20	4.07	155.00	22.00	-30.10	-10.03	174.20

Figure 2: Google Sheets spreadsheet containing the Side and Carry data collected for a participant's 7-iron through hybrid clubs. These measurements were exported from TrackMan and used in the simulation.

The Side and Carry distances for each player's clubs were imported into the Desmos graphing calculator, with Side on the x-axis and Carry on the y-axis. Desmos is a graphing calculator that allows users to create tables, visualize data by plotting points, and import images. Both variables were measured in yards. For each club, we calculated the average Side and Carry values and plotted them in a different colour from the 30 individual data points.

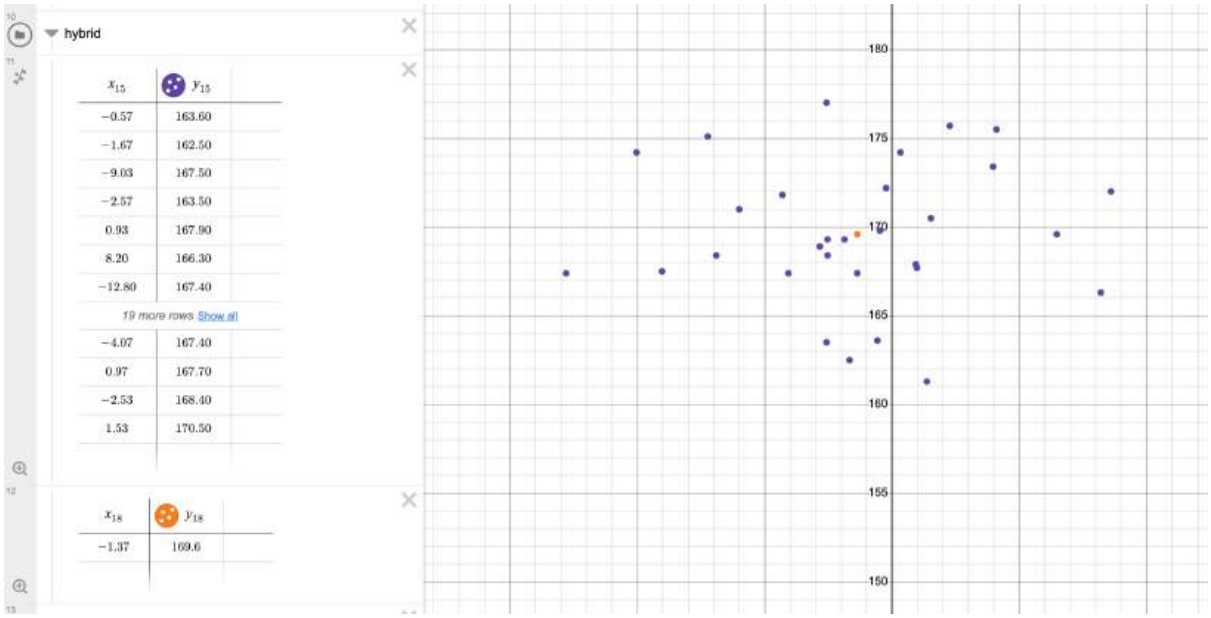


Figure 3: Scatter plot of the Side and Carry distances for a participant's 30 hybrid shots. The purple points represent the recorded shots, and the orange point indicates the mean shot location.

We ran the simulation from the White tees at Shaughnessy Golf and Country Club in Vancouver, which has a par of 74. Par is the number of strokes a skilled golfer is expected to require to complete the course under typical course and weather conditions. Based on the layout of the course, completing a round from the White tees typically requires around 38 full-swing shots, with the remaining strokes being mostly chips and putts. Therefore, the simulation assumed 38 full-swing shots per round. Following the 70-30 Rule, we chose 11 of 38 shots (approximately 28.9%) to be played aggressively, based on a game plan we constructed from the course layout. We made these choices ourselves, using our own experience playing competitive golf, rather than asking the participants to choose. For comparison, we created two additional strategies in which 20% and 40% of the full-swing shots were aggressive. We then compared the three plans to identify the holes played differently under each one. We highlighted these holes and simulated them using Google Earth. In total, 7 holes (13 shots) were simulated for each player. Because the other 11 holes were played the same across strategies, we did not include them in the comparison.

In this study, a shot was classified as *aggressive* if it met one of the following conditions:

1) Fairway target: The longest conventional club available was selected (excluding driver off the deck), with the center of the fairway as the target.
2) Green target: The club appropriate for the required distance was selected, with the pin as the target.

A shot was classified as *conservative* according to the following conditions:

1) Fairway target: When a larger or safer landing area was available farther from the green, a shorter club was selected, provided that doing so would increase the club required for

the subsequent shot by no more than two clubs and would not require a transition from an iron to a hybrid or a hybrid to a wood. If these conditions were not met, the longest conventional club available (excluding the driver off the deck) was selected. In either case, the target was the center of the fairway.

2) Green target: The club appropriate for the required distance was selected, with the center of the green as the target.

	chot 1	chot 2	chot 3
Hole 1	a	a	c
Hole 2	c	a	
Hole 3	c		
Hole 4	c	c	
Hole 5	c	a	c
Hole 6	c	a	
Hole 7	c	c	c
Hole 8	c		
Hole 9	a	c	
Hole 10	c	c	
Hole 11	c	c	a
Hole 12	c		
Hole 13	c	c	
Hole 14	c	a	
Hole 15	c	a	c
Hole 16	c	a	
Hole 17	a		
Hole 18	c	c	

Figure 4.1: Game plan following the 70-30 strategy, with approximately 30% of full-swing shots designated as aggressive.

	chot 1	chot 2	chot 3
Hole 1	c	a	c
Hole 2	c	a	
Hole 3	c		
Hole 4	c	c	
Hole 5	c	c	c
Hole 6	c	a	
Hole 7	c	c	c
Hole 8	c		
Hole 9	c	c	
Hole 10	c	c	
Hole 11	c	c	a
Hole 12	c		
Hole 13	c	c	
Hole 14	c	a	
Hole 15	c	a	c
Hole 16	c	a	
Hole 17	c		
Hole 18	c	c	

Figure 4.2: Game plan following the 80-20 strategy, with approximately 20% of full-swing shots designated as aggressive.

	chot 1	chot 2	chot 3
Hole 1	a	a	c
Hole 2	c	a	
Hole 3	a		
Hole 4	c	c	
Hole 5	c	a	c
Hole 6	c	a	
Hole 7	c	a	c
Hole 8	a		
Hole 9	a	c	
Hole 10	c	c	
Hole 11	c	c	a
Hole 12	c		
Hole 13	c	c	
Hole 14	c	a	
Hole 15	c	a	c
Hole 16	c	a	
Hole 17	a		
Hole 18	c	c	

Figure 4.3: Game plan following the 60-40 strategy, with approximately 40% of full-swing shots designated as aggressive. Orange highlights indicate the holes that differ between the three game plans and were selected for simulation.

Figure 4: The predetermined game plans used in the simulation, illustrating the distribution of conservative (c) and aggressive (a) shots across the 18 holes. Orange highlights indicate the holes that differ between the three game plans and were selected for simulation.

Google Earth is a geospatial software program that provides detailed aerial views of landscapes. In this experiment, we used it to determine each shot's projected landing location and assess the surrounding terrain. For each shot, a line was drawn from the point of impact to the projected landing location (Line 1). For simplicity, we assumed the player aimed accurately. Although this assumption removes aim variability, the shot dispersion data were collected from actual shots and therefore already incorporate natural variation. Thus, the simulation retains some realistic variability despite assuming accurate aim. A second line, perpendicular to Line 1, was drawn across the width of the fairway or the green at the landing location (Line 2).

We used Line 2 and the following steps to scale the golf course image in Desmos. We captured and imported a north-oriented screenshot containing Line 2 and the surrounding area into Desmos. We then rotated the image by the heading angle of Line 1, as measured in Google Earth. We measured the distances from Line 1 to the left and right edges of the fairway or green

in Google Earth, and plotted the corresponding vertical lines in Desmos at their respective x-values, with Line 1 represented by x=0. Line 2's equation corresponded to the y-coordinate equal to the length of Line 1, and the imported image was scaled to align with these reference lines.

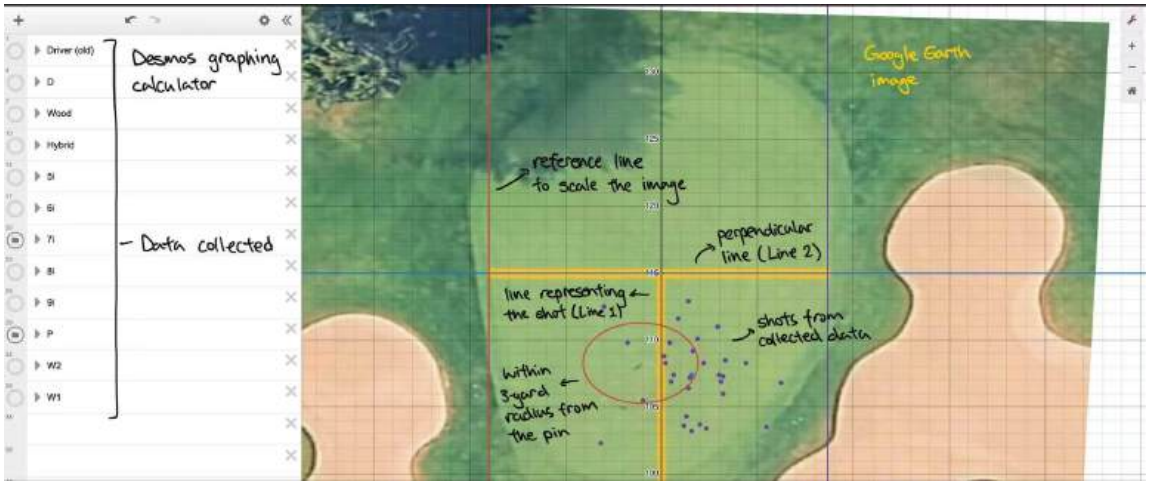


Figure 5: Annotated Desmos graph showcasing the shot distribution for participant Elsa's approach shot on Hole 9. The figure illustrates the alignment of the imported Google Earth image, the reference lines, and the recorded shot dispersion.

We then overlaid the shot dispersion for that club on the course image. This used the thirty shots we recorded during data collection. Each ball's landing location was classified into one of the following areas: within a 3-yard radius from the pin (y), on the green but outside a 3-yard radius from the pin (g), greenside fairway/fringe (gf), greenside rough (gr), greenside bunker (gb), fairway (f), first cut (fc), rough (r), bunker (b), trees/recovery position (t), hazard (h), and out of bounds (ob). We recorded each ball's location in a table. We repeated this procedure for all thirteen shots, with each shot represented by a separate column. We conducted the process for each of the three players under each of the three strategies, resulting in nine classification tables. For each location, we assigned a different Expected Points value, as shown in Table 1.

Location number	Location of the ball after a shot	Expected Points
0	Within 3 yards of pin (y)	0.1
1	Green, outside 3 yards (g)	0.7
2	Fringe (gf)	1
3	Greenside rough (gr)	1.5
4	Greenside bunker (gb)	2
5	Fairway (f)	2.5
6	First cut (fc)	3

7	Rough (r)	3.5
8	Fairway bunker (b)	4.5
9	Trees/recovery position (t)	6
10	Hazard (h)	7.5
11	Out of bounds (ob)	9.5

Table 1: Point values assigned to each possible ball landing location used in the Expected Points calculation.

The point values show the quality of the ball's resulting position, with lower point values corresponding to better positions. Positions closer to the pin are better, so a ball within 3 yards of the pin gets the lowest value and a ball out of bounds gets the highest. We based the scale on the Strokes Gained framework developed by Mark Broadie, a professor at Columbia Business School (Broadie). Strokes Gained compares how many strokes a golfer is expected to need to finish the hole before and after a shot. These expectations are estimated from historical PGA Tour data and account for factors such as lie type and proximity to the hole. As a result, our point scale is not evenly spaced and penalized based on lie type and distance to the hole.

Using RStudio, a program for writing and running R code, we simulated 100 rounds for each player under each strategy, using the data recorded in classification tables as described above. First, we created an R function to assign the point value to each landing location. Next, for each simulated round, using the sample() function in R, we picked one of the thirty recorded balls in a column of the classification table at random for each of the thirteen shots. We summed up the points received in each of the thirteen shots to get the Expected Points for the round. This procedure was nested in a for loop to repeat 100 times for each combination of the three players and three strategies, producing 900 Expected Points values in total. We then organized the 900 Expected Points values in a two-way table for statistical analysis, with the three strategies in separate columns and the 100 simulated rounds for each player in rows.

Lastly, we analyzed the 900 Expected Points table using a two-way ANOVA with replication. ANOVA is a test that determines whether the means of three or more groups differ by more than one would expect from chance alone. Our two factors were player (Players 1, 2, and 3) and strategy (80-20, 70-30, and 60-40), with 100 simulated rounds providing replicated Expected Points for each player-strategy combination. This allowed the test to measure the effect of each factor and determine whether they worked together. We used a significance level of $\alpha = 0.05$ throughout.

Because the ANOVA showed a significant difference among the three strategies ($p < 0.05$), we then ran two-sample t-tests that assume unequal variances (Welch's t-tests) to identify which strategies differed. We chose Welch's t-test because the group variances were unequal. We used two-tailed probabilities because we tested whether mean Expected Points differed between strategies, rather than predicting that one strategy would produce a specifically higher or lower mean. We made three comparisons: 80-20 versus 70-30, 70-30 versus 60-40, and 80-20 versus 60-40. Making several comparisons raises the chance of a Type I error, so we applied a

Bonferroni correction by dividing the original significance level of 0.05 by three (because of three comparisons). This gave an adjusted significance level of $\alpha = 0.0167$.

The ANOVA also showed a significant interaction between player and strategy ($p < 0.05$). This suggests that the effect of strategy differed among players, so we conducted the same three comparisons separately for each player. This resulted in nine t-tests. We used the same adjusted α level of 0.0167 for the three comparisons within each player.

Results

Strategies	80-20	70-30	60-40	Total
Player 1				
Average	21.09	16.21	16.52	17.94
Variance	4.60	2.69	3.61	8.60
Player 2				
Average	18.345	18.48	16.85	17.89
Variance	0.62	1.36	3.22	2.27
Player 3				
Average	21.13	23.3	24.17	22.87
Variance	4.91	6.27	8.06	8.01
All Players				
Average	20.19	19.33	19.18	
Variance	5.06	12.19	17.44	

Table 2: Summary statistics (mean and variance) of Expected Points for each player and game plan. Each cell is based on 100 simulated rounds.

Source of Variation	SS	df	MS	F	P-value
Sample (player)	4902.50	2	2451.25	624.25	3.29E-170
Columns (strategy)	177.60	2	88.80	22.61	2.63E-10
Interaction	1969.50	4	492.38	125.39	6.44E-85
Within	3498.72	891	3.93		
Total	10548.32	899			

Table 3: Results of the two-factor ANOVA with replication on Expected Points.

Player	Comparison	df	t Stat	P(T<=t) two-tail
1	80-20 vs. 70-30	185	18.07	1.04E-42
1	70-30 vs. 60-40	194	-1.24	0.218
1	80-20 vs. 60-40	195	15.95	3.33E-37
2	80-20 vs. 70-30	174	-0.96	0.338
2	70-30 vs. 60-40	170	7.62	1.71E-12
2	80-20 vs. 60-40	136	7.63	3.69E-12
3	80-20 vs. 70-30	195	-6.49	6.93E-10
3	70-30 vs. 60-40	195	-2.30	2.26E-2
3	80-20 vs. 60-40	187	-8.44	8.49E-15

Table 4: Results of the t-tests (two-sample assuming unequal variances) comparing Expected Points between game plans, within each player. Bonferroni-adjusted $\alpha = 0.0167$.

Table 2 presents the summary statistics for the Expected Points generated for each player under the three game plans. The 70-30 strategy gave the lowest average for Player 1 (16.21), the 60-40 strategy gave the lowest average for Player 2 (16.85), and the 80-20 strategy gave the lowest average for Player 3 (21.13). No single strategy was the lowest for all three players. Across all three players, the mean Expected Points were 20.19, 19.33, and 19.18 for the 80-20, 70-30, and 60-40 strategies, respectively. Table 3 shows the results of the two-way ANOVA.

The ANOVA revealed a statistically significant effect of player ($F = 624.25$, $p < 0.05$), suggesting that mean Expected Points differed across the three players. Strategy also had a significant effect ($F = 22.61$, $p < 0.05$), indicating that mean Expected Points differed across the 80-20, 70-30, and 60-40 strategies. Additionally, we observed a statistically significant interaction between player and strategy ($F = 125.39$, $p < 0.05$). This indicates that the effect of strategy differed across players. Table 4 shows the pairwise comparisons within each player.

For Player 1, the 80-20 strategy differed significantly from the 70-30 and 60-40 strategies. The 70-30 and 60-40 strategies did not differ significantly ($p = 0.218$). Although the 70-30 strategy gave Player 1 the lowest mean Expected Points (16.21), it did not differ significantly from 60-40. Therefore, there is insufficient evidence to conclude that the 70-30 strategy was better.

For Player 2, the 60-40 strategy was significantly different from the 70-30 strategy and from the 80-20 strategy. No significant difference was found between the 70-30 and 80-20 strategies. The 60-40 strategy produced the lowest mean Expected Points for Player 2 (16.85), and it differed significantly from the other two strategies.

For Player 3, the 70-30 strategy was significantly different from the 80-20 strategy, and so was the 60-40 strategy. The 70-30 and 60-40 strategies were not significantly different from

each other ($p = 0.023$, which is above the adjusted threshold of 0.0167). The 80-20 strategy gave Player 3 the lowest mean Expected Points (21.13), and it differed significantly from both other strategies.

Discussion

Previous studies have suggested that course-management strategy influences golf performance. The present study provides further evidence that supports this claim. The ANOVA revealed a statistically significant effect of strategy on Expected Points, indicating that mean Expected Points differed among the three strategies.

However, the pairwise comparisons did not demonstrate that one strategy consistently outperformed the others across all three players. For Player 1, the 70-30 strategy produced the lowest mean Expected Points, but its difference from the 60-40 strategy was not statistically significant. For Player 2, the 60-40 strategy produced significantly lower Expected Points than both the 70-30 and 80-20 strategies. For Player 3, the 80-20 strategy produced significantly lower Expected Points than both of the other strategies. This finding is noteworthy because it suggests that a course management strategy's effectiveness may depend on the individual player, rather than on a single optimal aggressiveness level for all golfers.

This study included only three participants, all junior female golfers. Consequently, the findings cannot be generalized to male golfers or golfers of different ages or skill levels, and because they are volunteers, the results apply only to golfers with similar characteristics. Furthermore, the simulations were based on TrackMan measurements collected under controlled indoor conditions and may not fully represent performance under varying outdoor course conditions. Additionally, only carry distances were considered for most clubs, and the amount of roll was assumed to be insignificant for simplicity. This assumption may affect the results, as roll remains significant in real-life settings. Finally, pressure during a tournament may also affect the results. Because the data were collected indoors during a practice range session, there was less pressure than there would be on the course. Golfers may make more errors on the course, but pressure can also improve their performance, so the effect is complicated and hard to measure in an experiment.

Because the Expected Points system was developed specifically for this study, the statistical results depend on the point values assigned to each shot outcome. The ordering of outcomes was determined based on their relative desirability in golf. For example, being 3 yards from the hole is generally better than being in the rough, which is better than being out of bounds. However, the numerical spacing between successive point values was based on subjective judgement. Alternative scoring systems that assigned larger or smaller differences between outcomes could yield different total Expected Points and, consequently, different statistical results. Future studies could run a sensitivity analysis by repeating the statistical analysis with alternative point scales to see whether the conclusions remain consistent. If similar conclusions emerge across multiple scoring systems, researchers can place greater confidence in the Expected Points system.

It is also worth noting that Expected Points depend on two factors at once: how good the player is, and which strategy they use. That is why we included player and strategy as separate factors, since a stronger golfer might score lower under every strategy. Therefore, the player effect and the strategy effect need to be examined separately. Future studies should also include a larger, more diverse sample of golfers to improve statistical power and assess whether the findings hold across populations. Additional studies could also examine golfers across different handicaps and ages, or assess whether the optimal strategy changes by hole or course difficulty.

Conclusion

This paper investigates the 70-30 Rule, first published by LuxGolf. The results showed statistically significant differences among course-management strategies, but did not identify a single best strategy across all players. This study provides an initial quantitative evaluation of the 70-30 Rule. By translating a qualitative strategy guideline into a measurable hypothesis, this research establishes a framework that future studies can refine and validate with larger datasets.

Although our results do not support the 70-30 split as optimal for all players, the rule highlights the role of strategic aggressiveness in game-plan design. The strategy associated with the lowest Expected Points varied among the three players, suggesting that the optimal balance between aggressive and conservative shots may be player-specific. Thus, rather than providing a fixed proportion, the findings suggest that consciously evaluating when and how often to play aggressively may be a useful framework for designing game plans. By providing a quantitative way to measure that balance, this study may help more junior golfers develop more informed and individualized game plans. More broadly, this research demonstrates how golf strategy can be examined through a quantitative lens. Building on previous studies in psychology, statistics, and decision theory, it contributes to the literature on how golfers make decisions under uncertainty.

This paper also suggests directions for future research. Because the data came from only three female junior golfers, a larger sample size is needed to provide more precise estimates and allow the findings to generalize more broadly. Additionally, as junior golf participation increases and data on junior golfers remain limited, this paper expands the statistical record of golf performance for this population. In particular, female golfers have traditionally been underrepresented in this sport, and most studies have relied on PGA databases or focused on male golfers. In the future, researchers should continue collecting data from female golfers to improve the representativeness and applicability of findings to women.

Ultimately, this study highlights the increasing role of statistics and quantitative analysis in understanding golf strategy. As more data becomes available, evidence-based decision-making can complement players' intuition and experience, leading to more informed course-management strategies.

Works Cited

Broadie, Mark. Assessing Golfer Performance on the PGA TOUR. Columbia University, 8 Apr. 2011, https://www.columbia.edu/~mnb2/broadie/Assets/strokes_gained_pga_broadie_20110408.pdf.

Broughton, Taylor S. Golf Brain: A Neuropsychological Study of Performance in Competition. 2021. George Fox University, doctoral dissertation.

Gnagy, E., et al. "An Exploration of Strategic Decision Making in Golf: Take a Chance, It's Worth the Risk." International Journal of Golf Science, vol. 4, no. 2, 2015, pp. 89–109, https://doi.org/10.1123/ijgs.2014-0016.

Hardt, Marie, and Dan Nettleton. "A Hierarchical Approach to Modeling Golf Hole Scores with Hardy Distributions." Journal of Quantitative Analysis in Sports, vol. 21, no. 2, 2025, pp. 97–115, https://doi.org/10.1515/jqas-2024-0036.

Khazaeli, M., and L. Javadpour. "Golf Club Selection with AI-Based Game Planning." Entropy, vol. 26, no. 9, 2024, article 800, https://doi.org/10.3390/e26090800.

Ko, S. A Simulation Model to Analyze the Impact of Golf Skills and a Scenario-Based Approach to Options Portfolio Optimization. 2012. Columbia University, doctoral dissertation.

LuxGolf. "What Is the 70/30 Rule in Golf?" LuxGolf, 2025, https://luxgolf.co.uk/70-30-rule-in-golf/.

Mack, Andrew B. A Model of Score Minimization and Rational Strategic Behavior in Golf. Harvard University OpenScholar, 2015.

McFall, Todd, and Kurt W. Rotthoff. "Risk-Taking Dynamics in Tournaments: Evidence from Professional Golf." Journal of Economic Behavior & Organization, vol. 179, 2020, pp. 378–394, https://doi.org/10.1016/j.jebo.2020.09.015.

National Golf Foundation. "Golf Industry Facts." National Golf Foundation, https://www.ngf.org/the-clubhouse/golf-industry-research/. Accessed 27 Sept. 2026.

Ozbeklik, Serkan, and Janet K. Smith. "Risk Taking in Competition: Evidence from Match Play Golf Tournaments." Journal of Corporate Finance, vol. 44, 2017, pp. 506–523, https://doi.org/10.1016/j.jcorpfin.2014.05.003.

Percy, David F. "Strategy Selection and Outcome Prediction in Sport Using Dynamic Learning for Stochastic Processes." Journal of the Operational Research Society, vol. 66, no. 11, 2015, pp. 1840–1849, https://doi.org/10.1057/jors.2014.137.

Pope, Devin G., and Maurice E. Schweitzer. "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes." American Economic Review, vol. 101, no. 1, 2011, pp. 129–157, https://doi.org/10.1257/aer.101.1.129.

Ramanayaka, Nilantha Dinesh, et al. "Heuristics in Sport: A Scoping Review." Psychology of Sport and Exercise, vol. 71, 2024, article 102589, https://doi.org/10.1016/j.psychsport.2023.102589.

Shapcott, S. Golfers' Decision Making on the Course: USGA Literature Review—Annotated Bibliography. United States Golf Association, 2021.

Stauffer, Guillaume, and Mathieu Guillot. “Golf Strategy Optimization and the ‘Drive for Show, Putt for Dough’ Adage.” Computational Statistics, vol. 40, no. 9, 2025, pp. 5381–5415, https://doi.org/10.1007/s00180-025-01659-6.

Stöckl, Michael, Peter F. Lamb, and Martin Lames. “A Model for Visualizing Difficulty in Golf and Subsequent Performance Rankings on the PGA Tour.” International Journal of Golf Science, vol. 1, no. 1, 2012, pp. 10–24, https://doi.org/10.1123/ijgs.1.1.10.

Wajge, Nishad, and Gautier Stauffer. “Should Sports Professionals Consider Their Adversary’s Strategy? A Case Study of Match Play in Golf.” Computational Statistics, vol. 40, 2025, pp. 2005–2029, https://doi.org/10.1007/s00180-024-01555-5.

Whitehead, Amy E. The Use of Think Aloud Protocol to Investigate Golfers Decision Making Processes. 2015. University of Central Lancashire, doctoral thesis.

Whitehead, Amy E., and Patricia C. Jackman. “Towards a Framework of Cognitive Processes during Competitive Golf Using the Think Aloud Method.” Psychology of Sport and Exercise, vol. 53, 2021, article 101869, https://doi.org/10.1016/j.psychsport.2020.101869.

Heteroscedastic Gaussian Process Regression for Out-of-Domain Neutrino-Nucleus Cross-Section Prediction

By Atharva Vignesh, Bhavya Darji, Yug Dhami, and Rohan Kumar

ABSTRACT

Uncertainties in neutrino interactions present a primary bottleneck for precision in particle physics because of the limited observable interactions. In an attempt to bridge this gap between the need for new discoveries and the limited available data due to the inherently low interaction cross-sections of neutrinos, our research team presents an uncertainty-aware Gaussian Process Regression model. Using data obtained from the MINERvA experiment, the data utilized in this study regarded the cross section of a muon neutrino's interactions with atomic nuclei of differing masses, 1D muon momenta, and uncertainty values; these nuclei were iron, water, lead, and carbon. Using a Held-Out Validation approach, the model was trained to learn the tendencies of muon neutrinos interacting with these nuclei, purposely leaving out Iron. This allowed us to evaluate the model's ability to generalize to an unseen target material. After testing using this approach,performance metrics, including RMSE and NLL, were calculated and visualized to evaluate model calibration; the graph pre-training produced a broad, nearly constant prediction interval and mean, whereas uncertainty calibration produced more dynamic predictions and narrower intervals. Analysis indicates, although RMSE increased from 0.1348 to 0.1575 after uncertainty calibration, this change does not necessarily indicate degraded model performance. Before training, the model produced a highly smooth prediction that remained nearly constant in terms of its cross-section ratio across multiple momentum values, resulting in a lower RMSE but limited response to variations in the data. After training however, the prediction became more dynamic and better reflected changes in momentum, while providing narrower and more responsive uncertainty intervals. This tradeoff is supported by the slight improvement in NLL, Negative Log-Likelihood, which decreased from −0.6365 to −0.6374. Since NLL evaluates both the predicted mean and its associated uncertainty, the improved NLL suggests that the trained model provides a better balance between predictive accuracy and uncertainty estimation. Ultimately, this framework clearly demonstrates how Gaussian Process Regression can reliably predict characteristics of neutrino interactions, offering an efficient surrogate model for advanced particle physics.

INTRODUCTION

Precision neutrino cross-section modeling faces three primary bottlenecks: First, experimental neutrino cross-section measurements are based upon limited, discrete data points across the periodic table. Most research facilities have access to very few elements; for instance, the data we utilized from MINERvA was based upon only four elements. As a result, it is extremely challenging to collect experimental data on additional materials. Consequently, the amount of available scientific phenomenon and discovery potential is reduced. This challenge is resolved with our model's ability to predict continuous cross-section ratios for varying atomic

masses, thereby obviating the necessity of individually prepared target materials. Furthermore, our Gaussian Process serves as a highly efficient surrogate model that rapidly substitutes for slow Monte Carlo simulations. Additionally, through direct incorporation of experimental uncertainty into a Bayesian framework, our Gaussian Process also provides calibration bounds for predicted values.

Precise modeling of neutrino interaction cross-sections (in particular muon neutrinos) on variable targets including carbon, water, iron, and lead will be essential to future experiments. Such experiments are typically conducted using equipment including MINERvA, DUNE, T2K etc. Since each of the above listed atomic nuclei produce different outcomes from interacting with a muon neutrino due to differing physical properties such as atomic number, a Gaussian process regression (GPR) model can aid in bridging the gap between the need for data covering many types of target atoms and those instances where there is limited observational data. Through the mapping of the GPR model over target atomic mass and one-dimensional muon momentum, the GPR model continuously learns about the data and has the capability to generate predictions for each property. This study utilizes data from the MINERvA experiment paper titled *Simultaneous measurement of muon neutrino νμ charged-current single π+ production in CH, C, H_2O, Fe, and Pb targets in MINERvA*, an experiment carried out by a well-known respected particle physics experiment.

Cross-section ratios are a primary way to quantify neutrino interactions. These ratios are a measure of the probability that a neutrino will interact with a certain molecule. This varies depending on numerous factors, and the two measurements that this research is based on are 1D muon momentum, and atomic number. Additionally, measuring cross-section ratios relative to hydrocarbon scintillators cancels out shared incoming neutrino flux uncertainties and detectors' systematic errors, leaving a clean measurement of atomic-mass-dependent nuclear modifications.

Furthermore, traditional Monte Carlo event generators (MCEGs) require the need for difficult, intense coding and computation, to be able to simulate particle motion and nuclear physics. By developing a data-driven surrogate model that explicitly accounts for point-wise experimental detector uncertainties, the framework we created delivers fast predictions and approximate uncertainty quantification across unobserved target materials. The rest of this paper is organized as follows: Section 2 details the MINERvA dataset and preprocessing protocols; Section 3 outlines the Gaussian Process Regression model architecture, covariance kernel selection, and Held-Out Validation; Section 4 evaluates model performance on held-out Iron (A=56), analyzing trade-offs between RMSE, Negative Log-Likelihood and Section 5 discusses physical implications and computational efficiency for future simulation pipelines.

DATASET & DATA PREPROCESSING

Data on high-energy neutrino interactions was obtained from the MINERvA experiment at Fermi National Accelerator Laboratory (FNAL). The MINERvA detector is a finely segmented, multi-target detector, which allows the simultaneous measurement of charged-current muon neutrino cross-sections on different nuclear targets under identical experimental beam

conditions. Data on cross-sections are recorded in the form of differential cross-section ratios taken relative to an internal hydrocarbon scintillator reference target. Such representation of cross-section measurements eliminates common uncertainties of the incoming neutrino flux spectrum, beam profile variations, and common systematic errors of the detector, leaving only pure atomic-mass-dependent nuclear effects and nuclear medium modifications across the observed momentum spectrum.

The dataset comprises four different targets with different atomic mass numbers, ranging from light nuclei (Carbon with atomic mass A = 12 and Water with atomic mass A = 18) to intermediate nuclei (Iron with atomic mass A = 56) and heavy nuclei (Lead with atomic mass A = 207). Each observation record consists of four main attributes: atomic mass number of the target nucleus (A), one-dimensional muon momentum (p_mu) spanning a continuous kinematic range from 0.5 to 16.0 GeV/c, cross-section ratio (y), and experimental measurement uncertainty derived from combined statistical and detector systematic measurement errors. Data on Iron with atomic mass A = 56 is reserved as an out-of-domain validation dataset, while Carbon, Water, and Lead are included in the training dataset in order to test the model's capability to extrapolate and interpolate predictions to other target materials across the periodic table without prior exposure.

Since input features vary strongly in their numerical scale (atomic mass numbers vary from 12 to 207, and muon momenta cover a continuous range in GeV/c), the use of unscaled inputs may result in distorted distances between samples in the input space and severely bias length-scale optimization in kernel-based machine learning methods. In order to avoid such distortion and ensure balanced feature weights, input features were standardized with Z-score normalization. It should be noted that the mean and standard deviation for feature scaling were computed only using observations from the training dataset in order to exclude any data leakage from the held-out validation dataset with Iron target material. These fixed parameters were then uniformly applied to transform both the validation set and continuous prediction grids.

Finally, experimental measurement uncertainties were carefully handled during data preparation. The uncertainties in the MINERvA dataset vary non-uniformly across momentum bins due to detector resolution limits and background subtractions. Rather than assuming constant noise across all data points, individual point-wise experimental uncertainties were extracted in their physical units and squared to form point-specific noise variances, preserving the heteroscedastic error structure of the dataset for model training and evaluation.

METHODOLOGY

Given that there are only a handful of measurements for each of the target materials, Gaussian Process Regression was used as the predictive model. More sophisticated models, e.g. neural networks, raise concerns about generalizability to inputs not seen during training, especially for sparse datasets. Moreover, the GPR model provides uncertainty estimates for predictions, allowing us to estimate the expected cross-section ratios for materials not included in the training.

Values were taken for muon momentum p_μ and cross-section ratio values, and their respective uncertainties, from *Simultaneous measurement of muon neutrino ν_μ charged-current single π+ production in CH, C, H_2O, Fe, and Pb targets in MINERvA*. The atomic mass A of the material for each data point is also noted. These are formatted as a dataset of inputs, $x = [A, p_\mu]$, target outputs (cross-section ratio), y, and target output uncertainties σ.

Carbon, Water, and Lead data are used as training data, while Iron data are held out and used as validation data, allowing us to test the model's predictive ability for an untrained target material. Because the magnitudes of each feature are not on a similar scale, the inputs are standardized as

$x' = (x - \mu_{train}) / \sigma_{train}$

where μtrain and σtrain are the mean and standard deviation computed using the training data only.

A simple GPR model with constant mean function is implemented using the GPyTorch framework. The covariance function is the product of individual Matérn 3/2 kernels for atomic mass and muon momentum and the output variance:

$$k(x, x') = \sigma_f^2 \cdot k_A(A, A') \cdot k_p(p_\mu, p_\mu')$$

$$k_A(A, A') = (1 + \sqrt{3}\cdot|A - A'| / \ell_A) \cdot \exp(-\sqrt{3}\cdot|A - A'| / \ell_A)$$

$$k_p(p_\mu, p_\mu') = (1 + \sqrt{3}\cdot|p_\mu - p_\mu'| / \ell_p) \cdot \exp(-\sqrt{3}\cdot|p_\mu - p_\mu'| / \ell_p)$$

σ_f^2 is the output variance, and ℓ_A and ℓ_p are length-scale hyperparameters for atomic mass and muon momentum, respectively. The model has lower and upper bounds set for its hyperparameters to avoid overfitting on only three materials.

The Matérn 3/2 kernel was specifically selected over others, such as the Radial Basis Function (RBF) or Matérn 5/2 kernels, because of how it suits sparse particle interaction data, and a lot of physical data in general. While the RBF kernel assumes infinitely differentiable, exceptionally smooth target functions, physical nuclear cross-sections often experience localized gradient changes across atomic mass numbers that can lead to overly stiff interpolation or numerical artifacts. The Matérn 3/2 kernel provides a once-differentiable prior that enforces continuous, smooth transitions between discrete nuclear targets without imposing unrealistically strict smoothness constraints.

The target output covariance

$$K_y = K(X, X) + \mathrm{diag}(\sigma_1^2, \ldots, \sigma_N^2)$$

incorporates point-wise measurement uncertainties directly onto the diagonal to implement a prescribed heteroscedastic GPR model. We use the uncertainties from the publication directly instead of learning the measurement noise during training, ensuring that more precise measurements will be weighted more strongly in the GPR's predictive distribution than less precise measurements. The model is trained by maximizing the marginal log-likelihood,

$$\log p(y \mid X) = -\tfrac{1}{2}(y - m)^T K_y^{-1}(y - m) - \tfrac{1}{2}\log|K_y| - (N/2)\log 2\pi$$

where m is the constant mean function, $K_y$is the target output covariance, N is the number of training examples, and $|K_y|$ is the determinant of K_y. Training proceeds using the Adam optimizer for 500 training steps.

Once the model has been trained, predictions are made using the held-out Iron data. Predictive performance is assessed using the following validation metrics: Root Mean Square Error (RMSE), Negative Log-Likelihood (NLL), and reduced χ^2. Predictions are also made for a grid of atomic mass and muon momentum values to examine interpolation between materials. Prediction time during inference is also reported.

EXPERIMENTS & RESULTS

To test the predictive ability and calibration of the uncertainty-aware Gaussian Process Regression model on unseen nuclear target materials, a Held-Out Validation scheme was implemented using the MINERvA dataset. Measurements from Carbon (atomic mass A = 12), Water (atomic mass A = 18), and Lead (atomic mass A = 207) were used to train the model and optimize its length-scale parameters, while measurements from Iron (A = 56) were held out completely and used as an out-of-domain validation dataset.

The model was created within the GPyTorch framework and standardized both atomic mass number and one-dimensional muon momentum using Z-score normalization based only on training-set statistics. The length-scales for the Matérn 3/2 covariance kernel, as well as the output variance scale, were then optimized by maximizing the marginal log-likelihood via the Adam optimizer over 500 iterations. Heteroscedastic measurement variances derived from published MINERvA point-wise total experimental uncertainties were included directly in the diagonal elements of the training covariance matrix.

Prior to model optimization, the uncalibrated model produced an overly smoothed, almost flat mean prediction across all muon momentum values. As indicated in Fig. 1, the baseline remains virtually flat across the entire momentum range from 0.5 to 16.0 GeV/c, with an excessively broad and nearly uniform prediction interval. The baseline model yields a lower point-wise root-mean-square error (RMSE) of 0.1348. However, this lower RMSE is misleading because it results solely from predicting a fixed global average near the center of the observations, completely missing the underlying momentum dependence and physically relevant variation in the cross-section ratio.

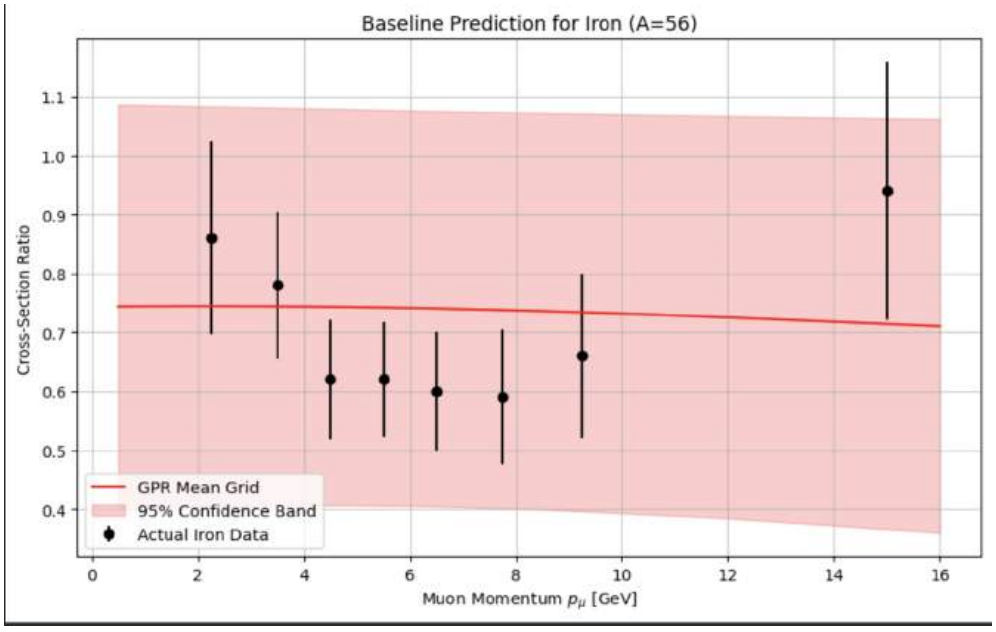


Figure 1: The initial baseline model prediction of a held out iron (with atomic mass A=56) is shown in Figure 1. In this figure, the red line shows the mean value from an un-calibrated grid-based GPR. The shaded red area surrounding this curve illustrates the wide 95% confidence band before any training occurred. The black circular markers represent individual MINERvA experimental data points for iron, which include their measured values along with their associated statistical uncertainties.

After optimizing its parameters using the Adam optimizer and explicitly incorporating experimental noise, the trained heteroscedastic model shows clear improvement in capturing physical features present in the experimental data over the baseline. As seen in Fig. 2, the trained model's mean prediction curve follows the non-linear curvature of the measured cross-section ratio. In particular, the mean prediction bends downward between 4.0 and 8.0 GeV/c to track the observed dip near 0.60. The model also separates uncertainty into a narrow latent function confidence interval (the dark blue shaded region) and a total confidence interval incorporating experimental noise (the light blue shaded region). This total confidence interval varies with momentum and widens in regions with greater experimental uncertainty, such as 15.0 GeV/c.

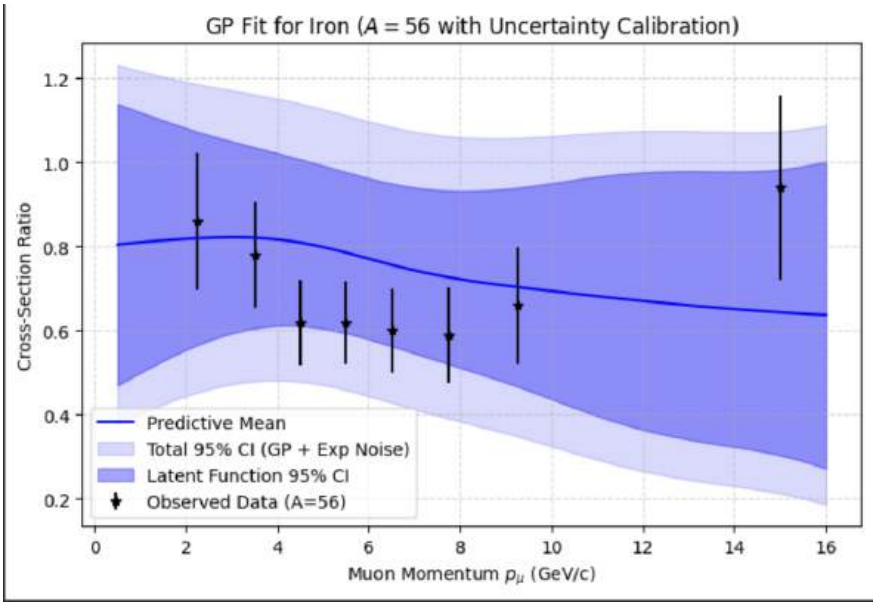


Figure 2: The Calibrated Gaussian Process Regression fit for the held-out Iron (atomic mass A = 56) is shown in this plot. The solid blue line represents the predictive mean of the model. The dark blue band represents a 95% Confidence Interval around the latent function that defines the underlying relationship. The light blue band represents a 95% Confidence Interval to account for experimental uncertainty. The black triangles with error bars are the observed MINERvA Iron data points.

Although post-training RMSE increases slightly from 0.1348 to 0.1575, this trade-off is justified by the strong qualitative and structural alignment between the trained model and the experimental data in Fig. 2. Rather than minimizing point-wise residuals with a flat, uninformative line, the trained model accurately tracks momentum-dependent dynamics while providing realistic two-tiered confidence intervals. This improved balance between physical trend fitting and uncertainty quantification is also supported by a small reduction in Negative Log-Likelihood (from -0.6365 to -0.6374) and a calibrated reduced chi-squared statistic of 0.8245.

Finally, to demonstrate continuous predictions beyond specific discrete target materials, the trained surrogate model was tested over a dense 2D grid covering atomic mass numbers from 12 to 207 and one-dimensional muon momenta from 0.5 to 16.0 GeV/c. As depicted in Fig. 3, the resulting interpolated surface demonstrates how cross-section ratios evolve smoothly across different atomic masses, providing a seamless bridge between the training targets (Carbon, Water, and Lead) and the held-out Iron target. Quantitatively, the interpolated surface captures physical nuclear shadowing trends: at low muon momenta below 4.0 GeV/c, lighter nuclei have cross-section ratios exceeding 0.825 (bright yellow region), whereas at higher momentum values, heavier nuclei such as Lead show cross-section ratio suppression down to approximately 0.345 (dark purple region). The smooth, continuous transition along the solid red line representing the held-out Iron nucleus (A = 56) confirms that the Matérn covariance kernel enforces realistic physical continuity without boundary artifacts or discontinuities.

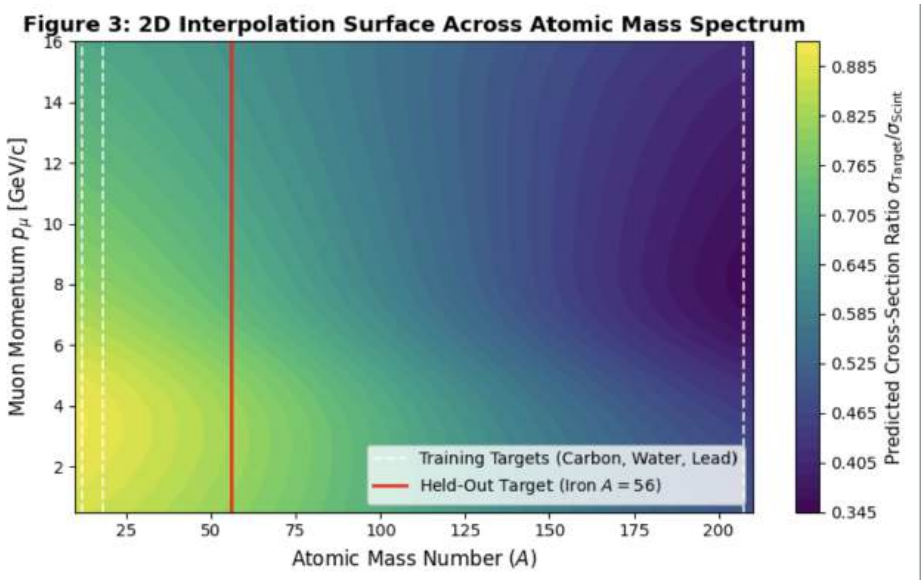


Figure 3:The predicted cross-section ratios in a two-dimensional representation as a function of atomic mass (from 12 to 207) and one-dimensional muon momentum (from .50 to 16.00 GeV/c) are shown above. The dashed vertical white lines indicate where the training data is located for Carbon, Water, and Lead; the solid red line at the right indicates where the validation data has been "held out" for Iron with an Atomic Mass of 56.

Furthermore, the model executed single-query predictions in 0.0279 milliseconds on CPU, offering an ultra-fast surrogate capable of replacing computationally expensive Monte Carlo event generators such as GENIE or GiBUU in high-energy physics pipelines.

DISCUSSION AND CONCLUSION

In this work, we presented a data-driven, uncertainty-aware Gaussian Process Regression framework designed to model continuous neutrino-nucleus cross-section scaling across atomic targets. Addressing the sparsity of physical target materials in experimental facilities, our approach demonstrates that non-parametric Bayesian models can successfully interpolate interaction cross-section ratios across continuous atomic mass numbers (A) and muon momentum (p_mu).

Using Held-Out Validation holding out Iron (A = 56), we demonstrated that incorporating known experimental error bars directly into the kernel covariance matrix yields highly calibrated prediction bounds (reduced chi-squared = 0.8245, NLL = -0.6374). Although uncalibrated models can yield deceptively lower RMSE values by predicting static averages, our analysis shows that joint likelihood metrics like NLL provide a far more reliable measure of physical model quality by balancing point accuracy with uncertainty quantification.

Ultimately, with its ability to process queries within 0.0279 milliseconds, it offers a fast surrogate model for the analysis of particle physics. In addition, by using the fast inference of a Gaussian Process to replace computationally intensive Monte Carlo event generation processes, researchers will be able to efficiently test nuclear systematic uncertainties for arbitrary choices of

target material that are part of their high-energy physics pipeline. In the future if the model is verified to produce accurate prediction and uncertainty assessments for other held-out target samples, it would allow researchers to focus measurement efforts on target masses and muon momenta where current experimental results have the least amount of constraining power. Additionally, the same approach could be used to estimate cross section ratios for neutrino experiments that involve materials scarcely characterized target materials.

Works Cited

Andreopoulos, C., et al. "The GENIE Neutrino Monte Carlo Generator." Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, vol. 614, no. 1, 2010, pp. 87-104.

Buss, O., et al. "Transport-Theoretical Description of Nuclear Reactions in the GiBUU Framework." Physics Reports, vol. 512, no. 1-2, 2012, pp. 1-124.

DUNE Collaboration. "Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report." Instruments, vol. 5, no. 4, 2021, p. 31.

GPyTorch Development Team. "Simple GP Regression with GPU Acceleration." GPyTorch Documentation, 2024, docs.gpytorch.ai/en/stable/examples/02_Scalable_Exact_GPs/Simple_GP_Regression_CUDA.html.

Bercellie, A., et al. “Simultaneous Measurement of Muon Neutrino ν_μ Charged-Current Single $\boldsymbol{\pi}^+$ Production in CH, C, H$_2$O, Fe, and Pb Targets in MINERvA.” 2209.07852v4, 16 Sept. 2022, arXiv, https://doi.org/10.1103/PhysRevLett.131.011801.

Rasmussen, C. E., and C. K. I. Williams. Gaussian Processes for Machine Learning. MIT Press, 2006.

Scikit-learn Developers. "Gaussian Process Regression with Noise Level Estimation." Scikit-learn Documentation, 2024, scikit-learn.org/stable/auto_examples/gaussian_process/plot_gpr_noisy_targets.html.

T2K Collaboration. "Characterization of Nuclear Effects in Muon-Neutrino Charged-Current Interactions on Hydrocarbon with the T2K Off-Axis Near Detector." Physical Review D, vol. 98, no. 3, 2018, p. 032003.