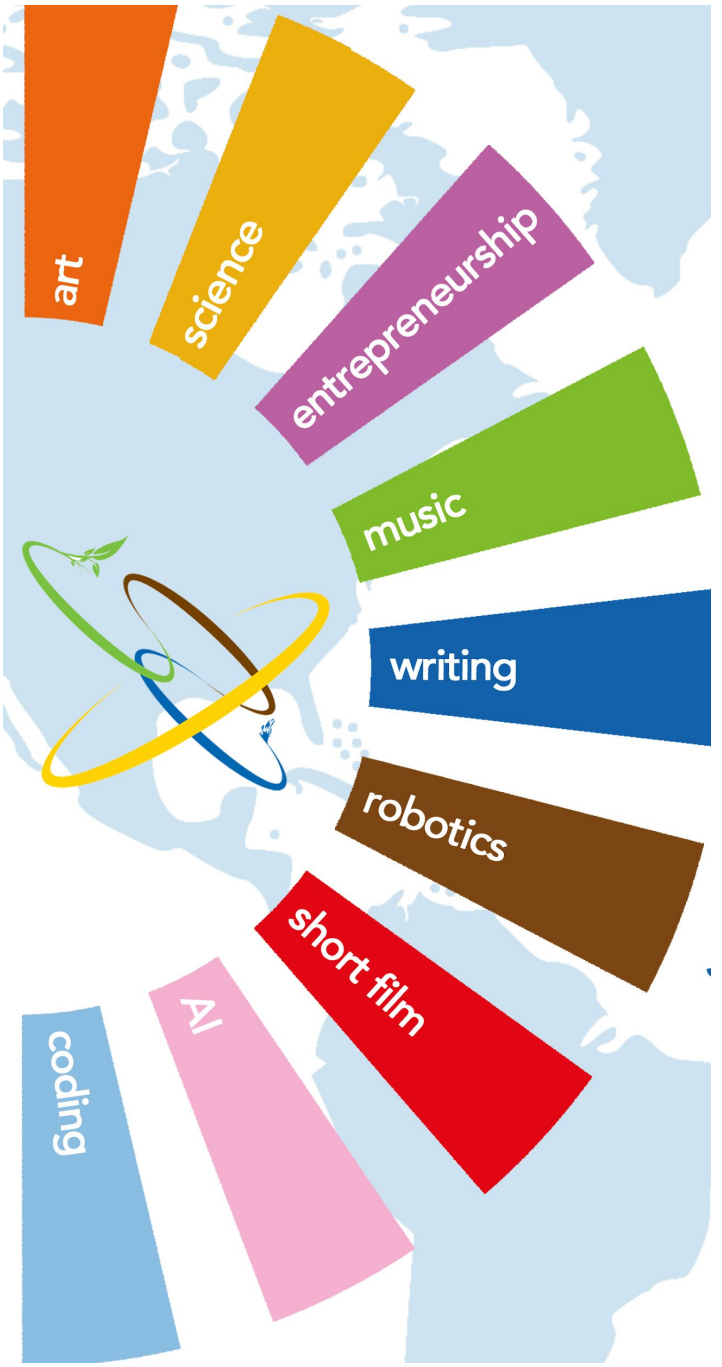
IJHSR International Journal of High School Research

October 2026 | Volume 8 | Issue 20
ijhsr.terrajournals.org
ISSN (Print) 2642-1046
ISSN (Online) 2642-1054

GENIUS OLYMPIAD

Let's build a better future together

International Project Fair focused on Sustainability and Environment for Grades 8-12



- Since 2011
- Hosted around 1400 participants in 2025 from 35+ states 70+ countries
- **Disciplines:** STEM, Coding, Robotics, AI, Speech, Entrepreneurship, Arts, Short Film, Music
- Applications start on December 1
- Application Deadline is March 1
- Finalists are announced by March 25
- Event is usually scheduled 2nd week of June
- Monday – Friday, includes a trip to Niagara Falls
- Hosted by large universities at Upstate New York
- Application Fee is $60/ project
- Participation Fee is $600/ person, w/ room and board
- Open buffet breakfast, lunch, and dinner
- Trip to Niagara Falls and boat tour is included
- **Instagram and Facebook @Geniusolympiad**
- For more information: **GENIUSOlympiad.org**
- Email: **info@geniusolympiad.org**

GENIUS Olympiad is organized by Terra Science and Education, a N.Y. based 501.c.3 non-profit organization dedicated for project-based learning

Table of Contents

October 2026| Volume 8| Issue 20

1 **Best Practices and Constraints in Installation and Maintenance of Solar and Wind Power: Lebanon Case Study**
Adam T. Zahran

14 **Exploring The Feasibility of Interstellar Propulsion: Assessing Sub-Century Travel to Alpha Centauri Under Modern Constraints**
Advait Y. Reddy

25 **Impacts of Early Intervention Plans on Cognitive and Socioemotional Development in Down Syndrome Patients**
Akanksha Mohanty

30 **The Vulnerability of Indigenous Rights In The Era of Global Extractivism: The Case of The Gateway Pacific Terminal**
Aliya Abdesheva

38 **Discovery of Small Molecules Targeting VEGF Using Virtual Screening as Potential Therapies for Vascular Malformations**
Andrew Wi

48 **The Effect of PIEB vs CEI for Quality of Parturient Care and Neonatal Outcomes: A Comparative Review**
Angelica C. Cheng

55 **Plasma Braking in Subcritical X-ray Pulsars: Coulomb Deceleration in Strongly Magnetized Atmospheres**
Chen Xingyu

63 **The Effect of Environmental Regulation Stringency on Productivity and Innovation: Testing the Porter Hypothesis**
Chris Z. Kumala

72 **How Can We Build an AI To Diagnose Textually Ambiguous Mental Health Disorders from Subjective Patient Reports?**
Dhairya Jain

77 **How Neurotransmitters Shape Depression**
Dhea Katkuri

84 **Machine Learning for Heatstroke Prediction: Unlocking Weather and Climate Signals**
Jingyao Zhu

91 **When Social Proof Becomes Demand: Price Resistance and Consumer Behavior in Viral Food Trends**
Elisa Lei

97 **The Electric Vehicle in the Automotive Industry - A Review of the Economic Barriers That Limit Its Large-Scale Adoption**
Kevi Sahara Karume

107 **A Comparative Laboratory Study of the Tensile Stiffness and Peel Strengths of Four Commercial Kinesiology Tapes**
Lester Yew Suan Lim

113 **Optimizing Ultrasonic Sensor Placement in Two-Sensor Systems for Reliable Navigation Using Dijkstra's Algorithm**
Mahika Kasarabada, Devin Carroll

117 **The Effect of the Reduction of the Cash Payment Limit on Compliance (VAT gap) in France**
Maria Sultanova

125 **The Application of Biomimetics to Maritime Transportation Through Antifouling and Hydrodynamics**
Max DeWeese

131 **Motorsport Brake Systems: Numerical Simulation and Analysis of Heat Distribution in Perforated Brake Discs**
Max B. Hong

145 **Special Functions: Gamma and Beta Functions - Theory and Applications**
Momo Kikuchi

152 **Developing Next-Generation dCas9-VP64 Systems for Temporal Activation Functionalities**
Rohaan Sanyal

158 **How Do Aluminum Formwork Systems Perform in High-rise Residential Buildings in India with Respect to Construction Efficiency and Reusability Under Real Site Conditions?**
Shivam Sharma

170 **Challenges and Possible Treatments for HER2-associated Breast Cancer**
Tien, Yu-Hsi

© 2026 Terra Science and Education

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Best Practices and Constraints in Installation and Maintenance of Solar and Wind Power: Lebanon Case Study

Adam T. Zahran

International College, Beirut, Lebanon; Adam.T.Zahran@gmail.com
Mentor: Professor Nazih Moubayed

ABSTRACT: This paper reviews best practices and constraints in the installation and maintenance of solar and wind power, laws and regulations that govern these systems in Lebanon, as well as an economic comparison of the cost of solar Photovoltaic (PV) systems compared to power from diesel generators. This paper examines the individual decentralized initiatives to shift from reliance on the national energy grid and diesel generators towards renewable energy sources as a way to achieve energy independence, triggered by the electricity situation in Lebanon, primarily after the country's 2019 financial crisis that led to widespread electricity interruptions and total collapse of the national electricity supply. This paper also highlights the importance of proper implementation of regulations that govern the installation of solar panels, which can be achieved through adequate and enforceable inspections.

KEYWORDS: Energy, Solar and Wind Power Generation, Installation and Maintenance, Best Practices, Lebanon Case Study.

Introduction

Lebanon has witnessed major electricity disruptions that have led to the reliance on diesel generators to power homes and businesses. The number of daily hours of electricity provided by Electricité du Liban (EdL) has witnessed a steady decrease to reach only 2-3 hours per day in mid 2021 onward, down from 11 hours in 2020. Diesel generators serve to fill the gap in electricity interruptions and provided in 2022 around 9 hours of electricity per day. [1]

The large reliance on diesel generators comes at a great cost, stemming from the economic, environmental, and health ramifications, because these generators are installed in highly populated areas.

Lebanon has strong potential for solar renewable energy with an average of 300 sunny days per year,[2] and potential for wind energy with average wind speeds in some areas reaching 7 m/s, particularly in the Akkar region and some coastal and mountain areas.[3]

Individuals and businesses have started tapping into this potential, especially in the absence of country-wide government initiatives to capitalize on renewable energy. Solar panel installation projects witnessed considerable growth between 2020 and 2023, with 15% of the country's total electricity generation in 2023 stemming from solar panels compared with only 2.2% in 2020.[4]

Although Lebanon issued in 2020 Decree No 6997 on the national standards relating to solar energy and photovoltaic systems and mandated the compliance with the standards that are issued by the Lebanese Standards Institution (LIBNOR), the current state of solar panel installations in Lebanon, as witnessed and documented firsthand following over 50 different rooftop site visits, reflects the lack of and disregard to proper installation guidelines, and the safety concerns that arise from lack of safeguards and inspection measures. Current solar panel installations also shed light on the lack of equity and fairness in accessibility to solar power.

Given Lebanon's current economic environment and the need to shift to a more reliable source of energy, Lebanon has great potential to capitalize on both solar and wind power, and can achieve this through the effective enforcement of its laws and regulations governing the usage and installation of solar panels, encouraging public-private partnerships to deploy solar and wind farm projects to help decrease the widespread chaos in the individual initiatives taking place and to reduce reliance on costly diesel generators, and employing AI tools for enhanced maintenance of both solar and wind power systems.

This paper reviews best practices for solar and wind energy installations and maintenance, and provides a look at Lebanon's potential and current challenges in implementing these renewable energy technologies. It also assesses the economic feasibility of transitioning to renewable energy as a pathway toward improved energy security through a cost comparison of diesel generators vs. solar and wind power.

Methodology

This paper combines a literature review with field observations and cost analysis. Primary literature and articles for this paper were retrieved from databases such as ScienceDirect and ResearchGate, and from institutional, governmental, and industry webpages. Field observations were conducted by the researcher, and digital photographs were taken to document findings. The cost data used in the economic comparison were derived primarily from IRENA and World Bank reports.

Discussion

Solar and Wind Energy Best Practices:

Solar Panels Installation:

The solar panel system is composed of several elements

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.1

and includes solar panels, a battery, an inverter, connecting wires, a grounding/earthing appliance, a mounting system, and any other components needed for proper functioning and harvesting of solar power. Figure 1 is a simple illustration of a residential solar panel system.

The most important step in adopting solar panels is a proper and safe installation method. Installation of solar panels on residential buildings, primarily on flat rooftops, requires careful planning, consideration, and assessment to evaluate the weight-bearing limit of these rooftops and to ensure that they offer sufficient structural stability and load-bearing capacity. The use of different anchoring techniques should not compromise the structural integrity of the rooftops and should take into consideration the impact of roof loading, weather elements, and wind resistance. A wind pressure coefficient based on the roof height and orientation needs to be accounted for when deciding on the optimal anchoring technique that will be used.[5]

One method of anchoring solar panels on rooftops is using a ballasted system, which is a noninvasive, nonpenetrative installation method that uses the heavy weight of concrete blocks as a way to secure the solar panels in place. These blocks are ideal for flat or low-slope rooftops (less than 5-degree slope), and are not fixed into the roof structure. The downside of using these concrete blocks is the potential cracking and corrosion that they can face due to their exposure to weather elements in Lebanon, including wind and rain.

The use of these concrete blocks necessitates a structural analysis of the rooftop to ensure sufficient structural capacity and support, and must account for high wind speeds, which would require additional anchoring in high wind areas, as these blocks might move and shift, which can damage the roof as well as expose the blocks to cracking.

The structural analysis needs to be done by a qualified team of professionals and needs to follow the proper structural codes.

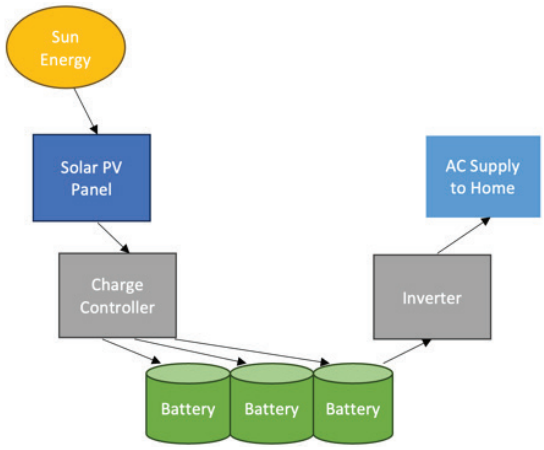


Figure 1: Solar panel system components.[6] The diagram illustrates the components needed for setting up a residential solar panel system, which includes the solar PV panel, which receives solar energy; the charge controller; the battery bank; and the inverter, which supplies AC power.

New technologies have emerged in the use of concrete blocks that add wind deflectors to reduce the risk of uplift. This aerodynamic feature aims at decreasing the amount of concrete used, in turn reducing the weight load on the building structure.[7] This new technology, if adopted in Lebanon, can help minimize the risk of adding heavy loads to aging rooftops.

Use of concrete blocks also needs to ensure equal weight distribution, as localized weight can have a long-term damaging effect on the structure of the roof. The use of these concrete blocks is low in cost as it is not highly labor-intensive.

Another method of anchoring solar panels is mechanically penetrating the roof and firmly securing the solar panels into the roof structure, which anchors them in place and offers a defense against strong winds in high wind areas. Depending on the roofing material used, mechanical anchoring requires different attachment techniques and waterproofing, as improper sealing can risk leaking into the roof structure through the holes in the roof. Special care and consideration need to be given to side building tiling when anchoring solar panels to the edge of the rooftop, as these tiles might come loose and require refixing.

The selection of material used in the mounting structure depends on the environmental conditions. Coastal areas in Lebanon with high humidity and salinity expose the structure to corrosion and rusting. While a steel structure might be considered a strong, cost-effective base, aluminum offers a lightweight alternative that has a high strength-to-weight ratio, given that its density is a third that of steel,[8] and is resistant to corrosion, rust, and other environmental factors. Use of aluminum would therefore give the structure a longer life cycle and would bring down the total cost of maintenance and replacement.

The mounting material and technique play an important role as well in ensuring that the entire solar panel structure is well anchored and that it can withstand high winds and not compromise the safety and security of passersby.

The efficiency of solar panels varies depending on the type of solar panel material used. Monocrystalline panels have an efficiency of 15-20% and use pure silicon crystals, whereas polycrystalline panels have an efficiency of 13-16% and use melted silicon crystals. In addition, monocrystalline solar cells have a lower temperature coefficient than polycrystalline cells, meaning that when the temperature rises, monocrystalline cells have a lower loss of efficiency compared to polycrystalline cells. As a result, monocrystalline cells are a better fit to serve hotter regions. The downside to monocrystalline cells is their high cost due to the price of the pure silicon crystals used in their production.[9] Given Lebanon's climate and the restrictive space on rooftops, monocrystalline cells prove to be the better option to maximize efficiency and minimize energy loss.

Monocrystalline cells are darker in color with round edges, which creates empty spacing between solar cells and would result in larger panels, whereas polycrystalline cells are blue in color with straight edges, as shown in Figure 2.

Use of high-quality materials during installation helps prolong the life of the panel system. The best approach to selecting batteries and inverters would be to prioritize using the same brand if feasible for easier installation and communication between components. Proper ventilation for batteries and inverters is crucial, as well as the use of DC wiring with proper grounding and cable protection against UV and moisture

damage. These measures help prevent fires or electrocution during maintenance and installation.

The physical placement of solar panels should aim for space optimization on rooftops to maximize energy generation, taking into consideration the required access space for cleaning and maintenance. Solar panels should be installed in areas with maximum sunlight exposure, avoiding shaded areas, which might be a problem in densely populated areas in cities in Lebanon, and areas with poor urban design where tall buildings obscure sunlight from surrounding shorter ones.

Different types of battery options can be used for energy storage, with Lithium-Ion batteries being the most widely used option due to their long life span, low lifetime cost, and low fire risk. Lead-Acid batteries have a low initial cost, but require frequent maintenance and more ventilation, and have a shorter cycle life. Large-scale installations rely primarily on flow batteries. A new emerging technology replaces liquid electrolytes used in lithium batteries with solid materials such as ceramics or sulfides, increasing energy density and eliminating liquid leaks. These solid-state batteries remain expensive due to the cost of materials. Table 1 summarizes the efficiency, energy density, cost, and cycle life of different types of batteries.

Table 1: AC power. Batteries are used in the storage of solar energy.[10]

Type of Battery	Efficiency	Energy Density	Cost $/kWh	Cycle life
Lithium Ion Battery	90%-95%	150-250 Wh/kg	Moderate: 250-500	2,000-5,000 cycles
Lead-Acid Battery	75%-85%	30-50 Wh/kg	Low: 100-200	500-1,000 cycles
Flow Battery	75%-85%	15-30 Wh/L	High: 500-1000	10,000 cycles

The capacity of the system is dictated by the number of solar panels used and is measured in kWp (kW peak), which represents the maximum power output under ideal conditions. The capacity is calculated using the power rating of each panel multiplied by the number of panels used. The solar panels are installed in series (string) and/or in parallel (array).

The simplest way to calculate the number of panels needed is to determine the daily energy needs based on historical energy data and the number of peak sun hours in the specific location.

Given that the rooftop areas in Lebanon are shared among residents of the building, another method of calculating the number of panels would be to determine the number of panels a rooftop can handle based on the available usable roof space. This method accounts for total roof area, safety clearance from the edge for maintenance and access, and space for other components installed on roofs such as vents, satellite systems, and water tanks.

It is recommended that the solar panel system produce 20% more power than the average energy needs to serve as a buffer for factors like energy loss during conversion, external parameters including cloud cover, and any energy usage spikes.

In setting up the solar panels, the ideal direction of the system is always in the direction of the equator, meaning towards the south in the northern hemisphere and towards the north in the southern hemisphere. The parameters needed to calculate the optimal tilt include latitude, time, day, and orientation (north/south).

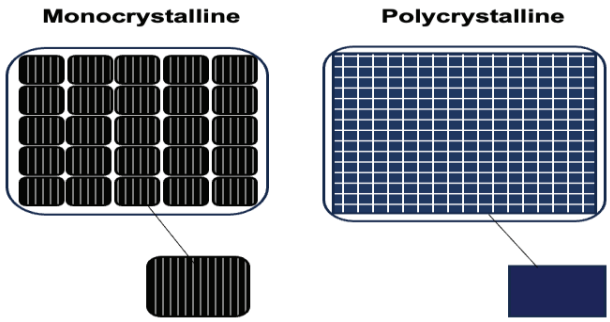


Figure 1: Different types of solar panels.[11] Illustration of monocrystalline panels, which are darker in color with round edges and formed from silicon, and polycrystalline panels, which are blue in color with straight edges and formed from fragments of silicon.

Various analytical formulas have been developed to calculate the optimal tilt position of the solar panels, and many have concluded that the annual optimal position is equal to the latitude of the location of the panels. The tilt of the panels can be adjusted monthly or seasonally to achieve higher performance, as some studies have concluded that adjusting the tilt 6 times per year led to 99.5% of performance and energy yield.[12]

The optimal solar panel tilt should be based on the solar elevation angle and should match the zenith angle. The Zenith angle is smaller in the summer months when the sun has a higher elevation and is larger in winter months when the sun is at a lower elevation. The azimuth angle represents the optimal horizontal direction of a solar panel (Figure 3).

A study done on Lebanon calculated the optimal annual tilt position to be averaging between 31.25 and 32.67 degrees for Beirut, between 30.91 and 32.16 degrees for Tyre, and between 31.58 and 33.15 degrees for Tripoli. The optimal tilt went up to the 60-degree range in winter months and down to the single-digit range in summer months.[13] The tilt should, in any case, exceed 10 degrees to create a slope for precipitation to flow, therefore limiting the accumulation of rain, snow, dirt, and dust on the panels.

The efficiency of solar panels is influenced by several external factors and parameters, including solar intensity, temperature, wind speed, cloud coverage, and humidity. A study done on the correlation between external parameters and the efficiency of the solar panels using linear regression concluded that a positive correlation exists between solar intensity and output power, whereas a negative correlation exists when it comes to temperature, cloud coverage, and humidity and their effect on output power.[14] With its high solar intensity, Lebanon can benefit from highly efficient solar energy if solar panel systems are correctly installed.

To increase the efficiency of systems and limit the effect of environmental factors, new technologies have emerged, including solar tracking systems that use a variable tracking device that sits on an axis to automatically change the tilt position of the panels. These tracking devices are costly compared to the fixed system that is manually adjusted to achieve the optimal tilt. This high initial cost, coupled with continuous upkeep, outweighs the added energy yield of the solar panel system.

DOI: 10.36838/IJHSR820.1

Solar Panels Maintenance:

Proper maintenance of solar panel systems is a crucial aspect of achieving high efficiency and maximizing the life cycle. Maintenance includes regular cleaning for dust, bird droppings, and pollutants, inspection and testing of batteries, inverters, and wiring.[15] Maintenance also includes repairing any roof leaks and checking for corrosion or cracks in the mounting structure.

Maintenance is a crucial element in Lebanon, given its high humidity, dust, pollution levels, as well as diesel smog and black carbon that accumulate on solar panels and severely impact their efficiency, the severity of which was witnessed and confirmed during rooftop site visits.

Cleaning panels starts with removing any particles and then spraying the panels with water to ensure proper cleaning. Regular cleaning and testing ensure optimal efficiency and limit energy loss. The optimal time for manual cleaning of solar panels is early morning or late afternoon when the system is not actively storing energy.

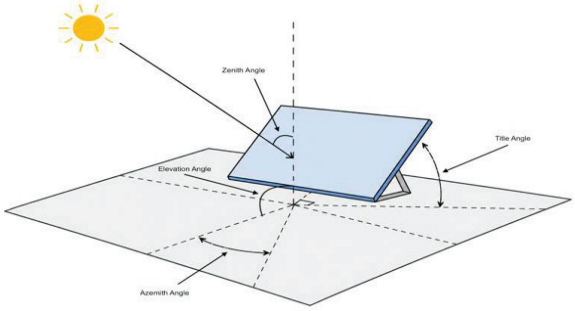


Figure 3: Optimal solar panel angles.[16] Relationship between the zenith angle and the optimal tilt of a solar panel, where the zenith angle is calculated relative to the vertical axis, and the tilt is measured relative to the horizontal axis.

Optimal maintenance of solar panels should rely on both passive and active cleaning. New innovative solutions are being deployed to optimize the efficiency of panels and reduce the need for manual intervention. Figure 4 lists the various types of mitigation techniques to decrease dust, pollutants, and buildup on solar panels, and to provide optimal conditions for solar panel performance.

Passive maintenance relies on environmental factors such as rain, wind, and gravity that help eliminate particles and decrease their adherence to the solar panels, and on the tilt of the solar panels, which helps these particles slide off the panels and limits their accumulation.

In contrast, active mitigation methods rely on intervening either through manual cleaning or through setting up automated systems to help with irrigation and cleaning. Manual cleaning should follow a specific schedule based on seasonal humidity, rainfall, and droppings in the area. AI tools can help with the scheduling of maintenance and cleaning of solar panels by gathering historical data on cleaning and maintenance of panels in specific areas, and proposing a cleaning and maintenance schedule.[17]

The incorporation of water-based systems is another method of active maintenance that does not require manual intervention. These systems are ideally used in dusty environments and areas with heavy bird droppings. These systems are installed on the panels and use a combination of automatic sprays and/or high-pressure water to remove dust and dislodge particles stuck on the surface of the panels. These water-based systems follow a preset daily or weekly schedule to optimize the efficiency of the solar panels. To reduce the amount of water used in the cleaning of solar panels, these water-based systems rely on recycling rainfall or water previously used for cleaning the panels, filtering it, and reusing it in the water-based system for panel cleaning.

Another method for active cleaning is the use of automated waterless techniques that involve the installation of small brushes that are deployed based on a specific schedule, which help remove dirt and dust particles and prevent them from sticking to panel arrays in humid climates. Another emerging technique involves the automated use of ultrasonic or acoustic waves that shake the dirt and dust particles off the panels.[18]

Use of remotely deployed robots in the cleaning of solar panels is gaining popularity, particularly for large solar panel systems, as they can assist manual cleaners and reduce the time required to complete the cleaning process.

Generative AI is shaping the future of solar panel maintenance by focusing on anticipatory and predictive vs. reactive maintenance to increase efficiency and reduce downtime, and is best used in large solar installations as opposed to residential areas. These systems are trained through the use of historical data and are fed with real-time data such as irradiance meters, thermal cameras, and inverter error logs to generate maintenance forecasting.[19] They also rely on satellite imagery and weather prediction data, including global horizontal irradiance, temperature, and humidity, to generate power forecasting for solar panels, and on drone thermal infrared footage and voltage readings to locate any potential problems and generate panel health reports and alerts for preventive maintenance.

Mitigation Methods							
Passive			Active				
Natural Mitigation	Coated Surfaces	Architectural and Design Solutions	Manual Cleaning	Water-Based Systems	Electro-mechanical Systems	Robotic Cleaning systems	Piezoelectric Systems
-Rain -Wind -Gravity	-Hydrophobic -Hydrophilic Emerging Technologies	-Tilt Optimization -Surface Microstructures		-Pressurized Water -Recycled Water -Surfactant Assisted Cleaning	-Brush Cleaning -Surface Acoustic Wave -Ultrasonic Cleaning -Electrodynamic Screens	-Autonomous -Semi Autonomous	

Figure 4: Various methods of cleaning solar panels.[20] The chart lists the various mitigation methods used in cleaning and maintenance of solar panels and includes both passive and active mitigation methods.

These systems can be beneficial for Lebanon, given the limited access space to conduct manual maintenance, and would therefore limit any risk of injury to the maintenance crew, and increase efficiency and energy output of the solar panel systems.

Wind Turbines:

While Lebanon can benefit greatly from a large-scale implementation of solar panels, wind proves to be another source of renewable energy that Lebanon can benefit from due to its high wind profile in a number of its areas. Wind turbines harvest energy from wind and convert the kinetic energy of

the wind into mechanical energy and ultimately into electrical energy.

The global reliance on wind energy has significantly increased in the last decade due to its efficiency and reliability as a source of sustainable energy. Many countries have expanded their use of wind energy, with Asia being at the forefront of this development. Figure 5 highlights the steady growth in installed wind power between 2011 and 2020, which reached over 700GW in global installed wind power.

Wind turbine installation takes place either offshore or onshore through what is referred to as wind farms. The taller the wind turbines, the higher the amount of energy generated, as these tall turbines access the stronger winds. Wind turbines include several components, starting with the blades and rotor, the hub, the nacelle, the generator, and the tower. Figure 6 provides a detailed view of the various components.

Wind turbines are classified as either horizontal or vertical based on their rotating axis. Horizontal wind turbines have their rotating axis parallel to the ground, whereas vertical turbines have their rotating axis perpendicular to the ground. Most wind turbines rely on the horizontal rotating axis, with the vertical ones primarily used in smaller spaces and residential areas. Figure 7 shows the difference between the two types.

The installation of onshore wind turbines requires careful planning and proper resources to ensure a smooth and safe installation and functioning of the wind turbines.

The site for installation of wind turbines is selected based on terrain and soil, wind quality and strength, as well as environmental factors. Onshore wind turbine installation requires logistical planning for the transportation of the large components of the wind turbines to the site, and resourcing large enough cranes for lifting the components to the desired height.

The tower is secured to the ground by the use of long bolts to anchor the structure into a cylindrical pedestal that is attached to a block of reinforced concrete. The Cylindrical pedestal has a diameter of 4 to 5 meters, typically, whereas the concrete base has a diameter of 15 to 20 meters and has a thickness of around 4 to 5 meters to allow anchoring deep into the ground. Figure 8 shows the different shapes of foundations that are used in securing wind turbines.

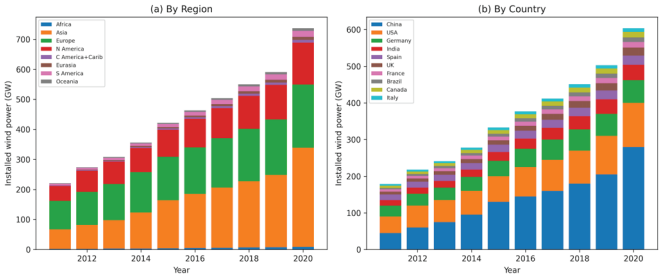


Figure 5: Global growth in wind turbines.[21] Charts illustrate the growth in installed wind power by continent and by country between 2011 and 2020, with Asia and China having the highest.

The bolts used in securing the tower to the cylindrical base are long anchor rods with lengths of around 6 meters that are placed into the reinforced foundation before the concrete sets. The tower of wind turbines is then attached and secured to these bolts, as shown in Figure 9.

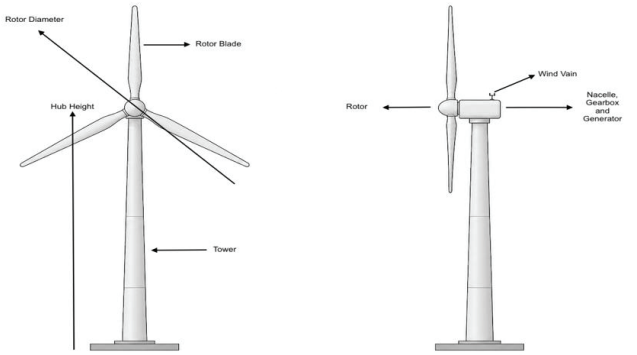


Figure 6: Components of wind turbines.[22] Representation of the main components of a wind turbine, which include the tower, rotor, and nacelle with gearbox and generator.

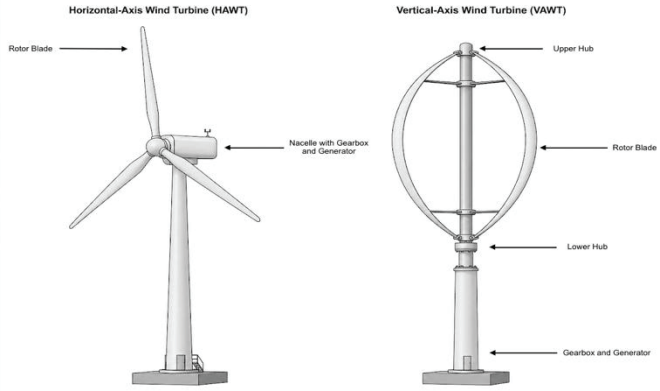


Figure 7: Horizontal vs. vertical axis of wind turbines. [23] Representation of a horizontal wind turbine with a rotating axis parallel to the ground versus the vertical wind turbine with a rotating axis perpendicular to the ground.

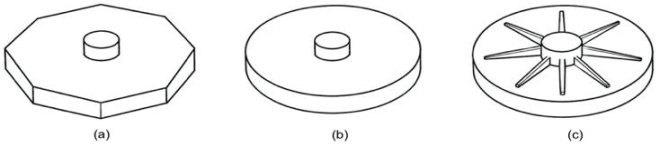


Figure 8: Typical wind turbine foundation configurations: (a) octagonal; (b) circular; (c) circular modular/segmented.[24]

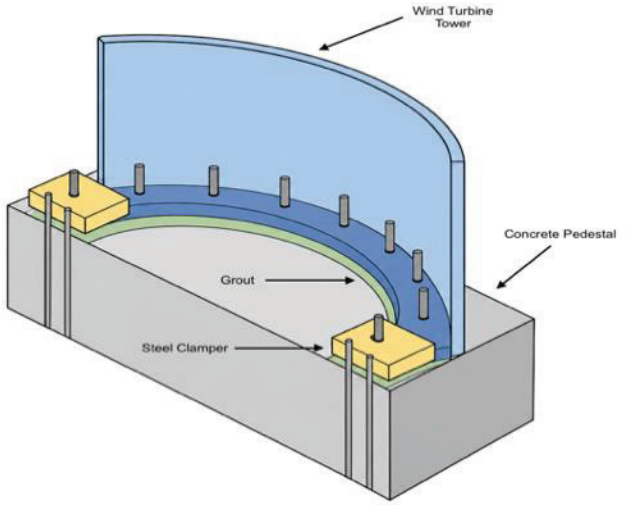


Figure 9: Representation of bolts used in the installation of wind turbines.[25] The illustration shows the long bolts used in securing the wind turbine tower into the concrete pedestal.

DOI: 10.36838/IJHSR820.1

The weight of the concrete foundation represents around 75% of the weight of the wind turbine system as a whole, and it offers stability to the system against wind and environmental factors. Figure 10 shows the massive size of the reinforced concrete.

The cost of operating and maintaining onshore wind turbines is fairly low and accounts for only 5% of their total cost,[26] where maintenance operations include servicing power and electrical wiring and systems, as well as maintenance of the main components, such as the blade and gearbox.

When it comes to undertaking an offshore wind farm initiative, the installation of turbines proves to be the most critical phase due to the cost as well as logistical and environmental constraints. Lebanon has yet to take on such an initiative, primarily due to its high cost and potential maritime border issues.

Installations of offshore wind turbines consist of different types of foundations, based on the size of the structure, the depth of the seabed, wind and wave intensity, as well as the cost and complexity of installation.

The majority of offshore wind turbines are bottom-fixed (Figure 11). These turbines can be fixed by either using a monopile foundation (a), which is suitable for depths of up to 40 meters, a gravity base foundation (b), suitable for water depths of up to 10 meters, or jacket-supported turbines, which are suitable for depths of up to 70 meters (c). For wind turbines installed in deeper water, floating wind turbines serve as the optimal solution and can be supported by anchors reaching the seabed (e). Lebanon's deep-sea floor requires it to use floating wind turbines, which are more complex and costly to install.

Installations are generally affected by weather and wind constraints, and require the use of vessels, ships, as well as cranes suitable for the transport and installation of the large tower and blades. Installation typically begins with laying the foundation, securing the tower, placing the nacelle and blades, and then securing the electrical wiring and connection.

Given Lebanon's location and high probability of seismic tremors, special consideration would need to be given in Lebanon to any potential seismic activity on the sea floor and its impact on the future of the wind turbines installed, as well as disturbing the marine ecosystems present in the selected location.

Several international standards and guidelines issued by DNV govern the installation of offshore wind turbines (Table 2). These industry standards include technical criteria and aim at ensuring the safety of the wind power structures, as well as the integrity of the structure.

New technologies in wind energy are emerging, some with the concept of including high-rise building- integrated wind turbines either on a horizontal or vertical axis, others with installing smaller appliance-sized wind turbines to generate an acceptable amount of energy, and are being used to power homes.[27] Another new technology in wind turbines revolves around harvesting energy from bladeless systems that rely on vibrations of cylindrical turbines. These new bladeless turbines offer great advantages relating to size, weight, and cost of harvesting energy, and can generate higher amounts of energy due to their ability to capture the vortex-induced vibrations from the wind flowing around a circular object.



Figure 10: Example wind turbine foundation.[25] The picture shows the size of the concrete foundation used in installing a wind turbine.

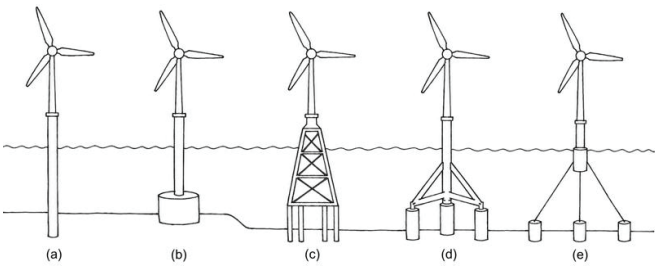


Figure 11: Types of offshore wind turbine installations.[28] Illustration shows different types of wind turbine foundations: (a) monopile foundation, (b) gravity base foundation, (c) jacket-supported foundation, (d) tripod foundation, (e) floating.

The new technological advances aim at reducing the cost of wind energy to achieve economies of scale and increase the total number of onshore and offshore wind turbines installed. Average costs of wind energy have been continuously decreasing and are expected to decline by up to 35% by 2035 and by up to 49% by 2050. These new cost levels are attributed to a larger turbine size that produces higher energy levels and lower operating costs. The following image portrays the expected increase in onshore and offshore wind turbine size and their corresponding power generation by 2035 (Figure 12).

Wind Turbine Maintenance:

Maintenance of offshore wind turbines poses great financial cost and risk, due in large part to the cost of acquiring the necessary equipment and logistical needs to carry out the maintenance operations, and the safety of staff involved in these operations. Operating and maintaining offshore wind turbines amounts to 23% of their total cost, making the maintenance of these wind turbines an essential phase of their life cycle. [26]

These maintenance operations can be categorized into three types of initiatives. Corrective maintenance takes place after a technical issue or failure is detected, which requires immediate remediation since it directly affects the level and amount of energy produced by the wind turbines. Preventive maintenance

relies on regular and frequent visits to prevent failures from happening in the first place. Opportunistic maintenance relies on performing maintenance whenever weather and other environmental and marine factors permit, regardless of the state of the wind turbines and the need for their maintenance.

Failure can take place in any part of the wind turbine system, with the tower and the blades accounting for the largest percentage of failures, as shown in Figure 13.

Table 2: Standards for installing offshore wind turbines.[29]

Standards	Title
DNVGL-ST-N001	Marine operations and marine warranty
ISO 29400:2015	Ships and marine technology – Offshore wind energy – Port and marine operations
DNV-RP-H103	Modeling and analysis of marine operations
DNVGL-ST-0378	Standard for offshore and platform lifting appliances
DNVGL-ST-0054	Transportation and installation of wind power plants
DNVGL-ST-0437	Loads and site conditions for wind turbine

Characteristics	Onshore Wind Turbines	Offshore
Wind Speed	7.5 m/s	9.5 m/s
Specific Power	231 W/m2	346 W/m2
Capacity	5.5 MW	17 MW
Rotor Diameter	174 m	250 m
Hub Height	130 m	151 m
Plant Size	-	900-1299 MW (fixed), 500-699 km (floating)
Distance to Shore	Land-based	70 km (fixed), 100 km (floating)
Water Depth	-	Up to 42 m (fixed), 100-199 m (floating)
Key Value Enhancements	Large rotors, storage hybrids, self-curtailement	Large rotors, balancing services, access to high-price and low-curtailment grid points
Main Barriers	Permitting, transportation, community concerns	Vessels, cranes, ports
Fixed-to-Floating Crossover	-	>60 m water depth
Use of Floating	-	11-25% of projects

Figure 12: Characteristics of wind turbines in 2035.[30] Comparison shows the expected characteristics of onshore and offshore wind turbines by 2035, including the increase in rotor diameter and hub height, with offshore wind turbines representing 11% to 25% of projects.

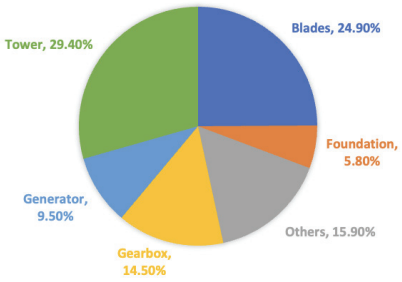


Figure 13: Percentage failures by system part.[31] The chart shows the failures wind turbines can face, with the highest failures affecting the tower and the blades.

Recent technological advances aim at maximizing the efficiency of maintenance operations and reducing costs by employing AI tools and monitoring systems to analyze environmental factors that might affect the offshore wind turbines and conduct fault diagnoses to ensure their efficiency and schedule maintenance accordingly. Current advances in AI models use data on mechanical stresses of the gearbox and the blade, including data on oil temperature and vibration of the gearbox, to generate early warning reports for parts replacement. AI models also rely on information on historical and current wind speed and wind flow to select optimal turbine placement and aerodynamic efficiency.

While advancements in the maintenance of wind turbines are paving the way for higher efficiency and more energy generation, high costs and financial constraints still pose implementation challenges, particularly in markets and regions like Lebanon, where the adoption of wind power is still underdeveloped. Nonetheless, these technological advances, if adopted, can provide long-term cost savings and can improve consistency and reliability.

Lebanon's Potential in Solar Power:

With many advancements in renewable energy driving costs of installation and maintenance down, Lebanon can greatly benefit from its location to implement these reliable sources of energy. Lebanon's Mediterranean climate allows it to benefit from year-round sunny weather, with a total of 300 sunny days, amounting to roughly 3000 hours of sunshine annually. This level of sunshine provides a solar irradiation level of approximately 1900 kWh/m2 per year, making it an ideal location to benefit from solar energy for both residential and commercial uses.[32] Figure 14, given by the World Atlas, represents the average photovoltaic power generation potential by area in Lebanon.

Laws Governing Solar Power in Lebanon:

Decree No 6997 on the national standards relating to solar energy and photovoltaic systems was issued in Lebanon on October 8, 2020, and mandates the use of the standards that are issued by the Lebanese Standards Institution (LIBNOR). These standards are based on international standards and cover areas from components of PV systems to installation and safety. The decree designates the Industrial Research Institutes (IRI) or any other internationally accredited laboratory as the party responsible for ensuring adherence and compliance with these standards (Figures 15 and 16).

In December 2023, the Lebanese Parliament ratified the Decentralized Renewable Energy Production Law No. 318/2023, which provides the legal framework for decentralized renewable energy production and sets the ground rules for the exchange of electricity (net-metering) with the national grid through Electricité du Liban (EDL) up to a maximum of 10 Mega-Watts (MW).[33] The implementation of this law is pending the establishment of the Electricity Regulatory Authority (ERA).

The concept of net metering allows owners of solar panels to share excess energy generated with EDL, with net metering calculated on a monthly basis by offsetting the amounts of energy provided to EDL with the amount of energy consumed.

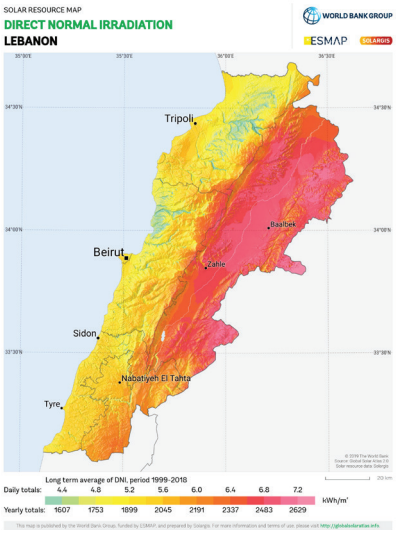
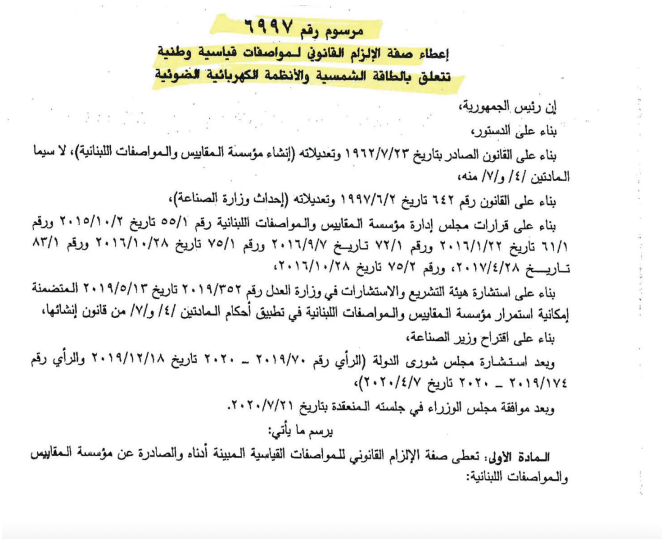


Figure 14: Photovoltaic potential by area in Lebanon.[34] The map shows the photovoltaic potential of different areas in Lebanon, with the highest potential in parts of the Beqaa Valley.

DOI: 10.36838/IJHSR820.1



مرسوم رقم ٦٩٩٧
إعطاء صفة الإلزام القانوني لمواصفات قياسية وطنية
تتعلق بالطاقة الشمسية والأنظمة الكهربائية الضوئية

إن رئيس الجمهورية،
بناء على الدستور،
بناء على القانون الصادر بتاريخ ١٩٦٢/٧/٢٣ وتعديلاته (إنشاء مؤسسة المقاييس والمواصفات اللبنانية)، لا سيما المادتين /٤/ و/٧/ منه،
بناء على القانون رقم ٦٤٢ تاريخ ١٩٩٧/٦/٢ وتعديلاته (إحداث وزارة الصناعة)،
بناء على قرارات مجلس إدارة مؤسسة المقاييس والمواصفات اللبنانية رقم ٥٥/١ تاريخ ٢٠١٥/١٠/٢ ورقم ٦١/١ تاريخ ٢٠١٦/١/٢٢ ورقم ٧٢/١ تاريخ ٢٠١٦/٩/٧ ورقم ٧٥/١ تاريخ ٢٠١٦/١٠/٢٨ ورقم ٨٣/١ تاريخ ٢٠١٧/٤/٢٨، ورقم ٧٥/٢ تاريخ ٢٠١٦/١٠/٢٨،
بناء على استشارة هيئة التشريع والاستشارات في وزارة العدل رقم ٢٠١٩/٣٥٢ تاريخ ٢٠١٩/٥/١٣ المتضمنة إمكانية استمرار مؤسسة المقاييس والمواصفات اللبنانية في تطبيق أحكام المادتين /٤/ و/٧/ من قانون إنشائها،
بناء على اقتراح وزير الصناعة،
وبعد استشارة مجلس شورى الدولة (الرأي رقم ٢٠١٩/٧٠ ـ ٢٠٢٠ تاريخ ٢٠١٩/١٢/١٨ والرأي رقم ٢٠١٩/١٧٤ ـ ٢٠٢٠ تاريخ ٢٠٢٠/٤/٧)،
وبعد موافقة مجلس الوزراء في جلسته المنعقدة بتاريخ ٢٠٢٠/٧/٢١،
يرسم ما يأتي:
المادة الاولى: تعطى صفة الإلزام القانوني للمواصفات القياسية المبينة أدناه والصادرة عن مؤسسة المقاييس والمواصفات اللبنانية:

Figure 15: Decree No 6997 (Arabic) – introduction.[35] The introduction section of the decree lays the legal obligation to comply with the standards issued by the Lebanese Standards Institution (LIBNOR).



الجريدة الرسمية ـ العدد ٣٩ ـ ٢٠٢٠/١٠/٨ ١٥٩٦

عنوان المواصفة	الرقم
Photovoltaic devices - Part 1: Measurement of photovoltaic current-voltage characteristics	NL EN 60904-1:2015
Thermal solar systems and components - Solar collectors- Part 1: General requirements	NL EN 12975-1:2015
Solar energy - Vocabulary	NL ISO 9488:2015
Solar Heating - Domestic Water Heating Systems - Part 2: Outdoor Test Methods for System Performance Characterization and Yearly Performance Prediction of Solar-only Systems	NL ISO 9459-2:2015
Solar Heating - Domestic Water Heating Systems - Part 5: System Performance Characterization by Means of Whole-System Tests and Computer Simulation	NL ISO 9459-5:2015
Balance –of system components for photovoltaic systems- Design qualification natural environments (IEC 62093:2005)	NL EN 62093:2017
Safety of power converters for use in photovoltaic power systems - Part 1: General requirements (IEC 62109-1:2010)	NL EN 62109-1:2016
Safety of power converters for use in photovoltaic power systems - Part 2: Particular requirements for inverters (IEC 62109-2:2011)	NL EN 62109-2:2016
Utility-interconnected photovoltaic inverters-Test procedure of islanding prevention measures (IEC 62116:2014)	NL EN 62116:2016
Photovoltaic systems - Power conditioners - Procedure for measuring efficiency. (IEC 61683:1999)	NL EN 61683:2016
Battery charge controllers for photovoltaic systems - Performance and functioning. (IEC 62509:2010)	NL EN 62509:2016
Solar photovoltaic energy systems - Terms, definitions and symbols	NL CLC/TS 61836:2015

المادة الثانية: تعتمد نتائج التحاليل والاختبارات الصادرة عن مختبرات معهد البحوث الصناعية أو اي مختبر آخر حاصل على الاعتماد الدولي للاختبارات المطلوبة، للتحقق من مطابقة المنتجات المعنية بهذا المرسوم للمواصفات المحددة في المادة الاولى أعلاه.

المادة الثالثة: تعتبر المنتجات المصنعة محليا الحاصلة على شارة المطابقة من مؤسسة المقاييس والمواصفات اللبنانية لإحدى المواصفات القياسية الواردة في المادة الأولى أعلاه، مطابقة حكما لمتطلبات المواصفة المعنية، مع مراعاة أحكام المادة الثانية اعلاه عندما تحدد المؤسسة حاجتها للقيام بذلك.

المادة الرابعة: تلغى جميع النصوص والأحكام المتعارضة مع أحكام هذا المرسوم.

المادة الخامسة: يعمل بهذا المرسوم فور نشره في الجريدة الرسمية ويبلغ حيث تدعو الحاجة.

بعبدا في ٢٩ ايلول ٢٠٢٠
الامضاء: ميشال عون

صدر عن رئيس الجمهورية
رئيس مجلس الوزراء
الامضاء: حسان دياب

وزير الصناعة
الامضاء: عماد حب الله

Figure 16: Decree No 6997- specifications.[35] This section includes the standards governing the PV solar devices, systems, terms, definitions, and symbols.

This law also allows for the exchange of electricity between residents of neighborhoods within a specific geographical area, as well as for municipalities to sell solar energy to their residents, without the need to go through EDL.

The transition to decentralized solar energy in Lebanon has emerged in response to the inability of the national electricity system to provide reliable power. The solar energy boom and transformation in Lebanon stems from an energy crisis that was intensified by the financial and energy system collapse of October 2019. However, the decentralized current state of adoption of solar technology and installations in Lebanon shows the amount of chaos, injustice, and lack of safety in the installation of solar systems on residential rooftops.

Many installations of solar panels in Lebanon are performed by unqualified individuals who do not account for wind load in high wind areas. The installation of solar panels in Lebanon for personal use of up to 1.5 Mega-Watts does not require approval from the Ministry of Energy and Water. The installation of the metal structures, however, is under the Directorate General of Urban Planning and requires the acquisition of the appropriate permits.[36]

A study done in 2022 by the Lebanese Center for Energy Conservation (LCEC) on the state of solar panel installation in Lebanon shows that in 2020, residential solar panel installation in Lebanon, in terms of kWp (kW peak), had increased by 93% year on year, with an investment of 26.6 million USD in the residential sector.[37] The most used mounting system type according to the study is the railed mount, followed by the foundation mount, and then the ballasted mounting system.

Site Visits and Observations:

Several site visits were conducted in the last quarter of 2025 in Beirut and neighboring areas to observe and document firsthand the current state of implementations of solar panels on rooftops. Sites were selected in areas with older buildings, newer buildings, and highly condensed areas. Observations focused on several criteria, including the mounting structure and foundation of the system, wiring, service clearance, corrosion, and dust accumulation. The observations did not address the amount of energy generated by the system, but focused primarily on installation. Digital photographs were taken of installations in 50 different sites. Photographs taken highlighted several issues regarding structural integrity, electrical safety, accessibility, corrosion, and the lack of cleaning of panels. Table 3 summarizes the main issues observed in solar installations and the frequency of these issues.

Table 3: Decree No 6997- specifications.[35] This section includes the standards governing the PV solar devices, systems, terms, definitions, and symbols.

Issue	Frequency
Structural Issues	17 out of 50 sites had structural issues
Electrical Safety	35 out of 50 sites had wiring exposed to weather elements
Service Clearance	32 out of 50 sites had limited access for maintenance
Corrosion	27 out of 50 sites had signs of rusting and corrosion
Dust Accumulation	41 out of 50 sited had dust accumulation

These field observations revealed the need to implement specific protocols with annual inspections to mitigate the risks, vulnerabilities, and hazards. Protocols need to address at a minimum the following needs to ensure safe and proper functioning of installed systems:

- Certified installers
- Load calculation protocols
- Approvals from structural engineers
- Anchoring and mounting protocols

- Electrical code for grounding and electrical testing
- Mandatory maintenance access
- Mandatory spacing from other utilities on the rooftop
- Use of anti-corrosion metal, such as aluminum, for the frames and mounting structure

A sample of these photographs is included in figures 17-30. Rust on the steel mounting structure, as shown in Figures 17 and 18, poses risks to the integrity of the system. This confirms aluminum as the better material of choice in installing solar panels in Lebanon due to its resistance to weather and environmental elements. Lack of space around the solar panel system, as seen in Figures 19, 20, and 21, renders the maintenance, cleaning, and inspection of solar panels both unsafe and ineffective. Wiring exposed to weather and environmental elements in Figure 22, compared to cable protection used in Figures 23 and 24. No adherence to regulations in the use of space on rooftops, as solar panels are installed over water tanks, as seen in Figure 25. Figure 26 shows the effect of high wind on a poorly installed solar panel, which raises concerns about the safety of passersby, an issue raised by YASA (Youth Association for Social Awareness).[38]

Given the current political and economic environment in Lebanon, and the lack of equity and fairness in the current state of solar panels in residential buildings, community solar builds on previous experiences of resorting to neighborhood public generators to power homes, and has emerged as a viable solution to improve access to electricity and decrease costs. By aggregating resources, participants in community solar can benefit from shared installation, making access to generated energy more affordable and efficient.

Access to empty pieces of land poses the main challenge in implementing community solar. Covering parking lots with solar panels can serve as one solution to the limited available space in cities. Figures 27-30 are taken from a current community solar park being built. Given the size of the project and the monetary investment allocated, these figures show more adherence to safety rules and regulations, primarily with no exposed wiring.



Figure 17: Rust on the steel structure of a solar panel. The image shows the negative impact of using a steel structure in an environment with high humidity.



Figure 18: Rust on the anchoring connection of a solar panel installation. The image shows the impact of rust on the integrity of the anchoring system.

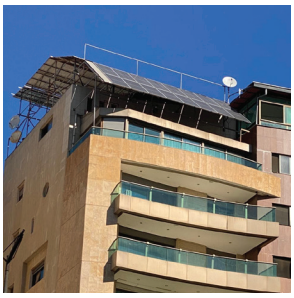


Figure 19: Solar panel installation. The image shows the installation of solar panels on rooftops with no spacing for maintenance work.



Figure 20: Solar panel installation on the edge of the rooftop. The image shows the installation of solar panels extending beyond the rooftop, which poses safety issues.



Figure 21: Anchoring of the solar panel on the edge of the building. Image shows the lack of planning and safety measures, as solar panels are anchored on the edge of the building, which allows for minimal space for maintenance.



Figure 22: Exposed wiring of a solar panel installation. Image shows the lack of safety measures as wiring used is exposed to weather elements.



Figure 23: Covered wiring of a solar panel installation. The image shows the use of covered wiring in the installation of solar panels.

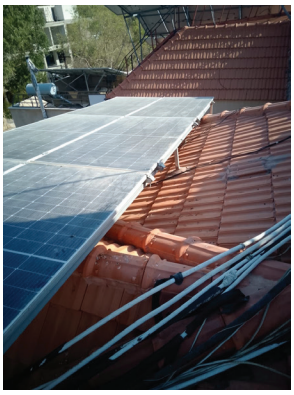


Figure 24: Wiring placement on rooftop installation. The image shows the placement of the wiring of solar panels.



Figure 25: Solar panel installation over water tanks. The image shows the lack of planning in space use on rooftops and the installation of solar panels and wiring over preexisting water tanks.



Figure 26: Effect of high wind on solar panels.[38] The image shows the collapse of solar panels on a road following a gust of high winds.

DOI: 10.36838/IJHSR820.1



Figure 27: Solar Park installation. The image shows an example of an installed solar park.



Figure 28: Site preparation for solar park. The image shows the extensive work needed for site preparation for the installation of solar panels in a solar park.



Figure 29: Anchoring is used in a solar park. The image shows the use of aluminum for the structure and anchoring of solar panels.



Figure 30: Battery installation in a solar park. The image shows an installation of a battery with covered wiring.

Lebanon's Potential in Wind Power:

Lebanon's wind power potential is substantial and can provide the country with the much-needed relief in terms of electricity shortages. Figure 31 shows the wind speed in different locations in Lebanon at an 80-meter height.

As part of its national strategy of increasing its reliance on renewable energy sources and reducing its greenhouse footprint, Lebanon has started to tap into its onshore wind farm potential by signing during the year 2018 three onshore wind farm projects located in North Lebanon in the Akkar region under the names Lebanon Wind Power (LWP), Hawa Akkar Wind Farm, and Sustainable Akkar Wind Farm, to install a total of 54 turbines generating a total of 219 megawatts (MW). These projects have yet to be completed and made fully operational.[39]

Once operational, these projects will face some challenges primarily due to Lebanon's, particularly Northern Lebanon's, location in the migratory route of hundreds of millions of protected birds flying from Africa to Europe in the Spring, and from Europe to Africa in the Fall. Wind turbines operate at the same height and corridor that birds fly into. Measures would need to be taken to protect the safe passage of these birds, such as shutting off or slowing down the wind turbines during the migratory period. AI tools can also be used that can monitor and detect birds flying into the wind turbine region, and signal automatic shutdown of the blades until safe passage is completed.

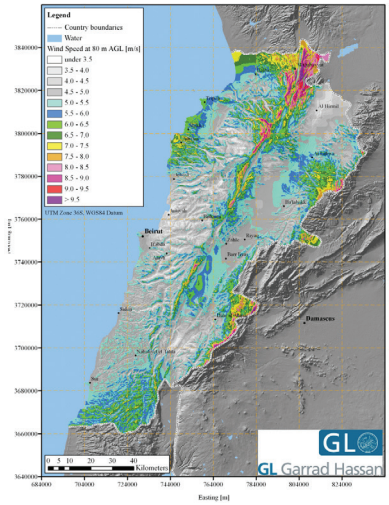


Figure 31: Wind speed in different areas in Lebanon.[40] The map shows the different wind speeds for different areas in Lebanon, with the highest wind speed in the Akkar region.

When it comes to offshore wind farms, Lebanon's offshore potential is largely affected by the high costs of such initiatives, its technical complexity, and the logistical implications of installing the foundational structures of wind turbines in the deep Mediterranean Sea.

Economic Comparison of Diesel Generators vs. Solar and Wind Power:

The cost of diesel per kWh has been increasing in the last few years and has reached a level of 55 cents per kWh in May 2026 due to the ongoing conflict in the region and its effect on fuel prices.

In contrast, technological advances have driven down the upfront costs of solar and wind power systems, with a 90% decrease in the cost of Photovoltaic systems (PV) between 2010 and 2023, putting these renewable energy sources at the lowest levelized cost of energy LCOE. The average LCOE in 2023 of wind power systems is the lowest at 0.03 USD/kWh, followed by solar power with 0.04 USD/kWh, and diesel generators with 0.44 USD/kWh. By combining the three sources of energy, where diesel generators are used to cover the gap in solar and wind energy produced, the LCOE of a hybrid system amounts to an average of 0.28 USD/kWh.[41]

The initial capital expenditures (CAPEX) for solar and wind energy are high compared to diesel generators. In return, diesel generators have high fuel and other operational expenditures (OPEX), which make up about 80% of their LCOE. Despite having higher initial costs than diesel generators, both solar and wind power are economically more beneficial over the lifetime of the system. In contrast, EdL's reforms aim to set the cost of electricity at $0.27/kWh. Table 4 provides a look at these costs for the different types of energy sources in Lebanon, and the LCOE after accounting for maintenance and fuel price at $0.93/L for diesel generators.

Table 4: Cost comparison of kWh per source and cost of the system to generate 10 kW in Lebanon.[4, 42, 43]

Type of Energy Sources	CAPEX (Materials & Installation)	Cost for 10kW	Fuel	OPEX	LCOE
Solar	~$1000/kW	~$10,000	None	Low- 1%-2% of CAPEX	$0.08/kWh
Wind	~$1,600/kW	~$16,000	None	Low-Moderate	$0.06/kWh
Diesel Generator	~$200/kW	~$2,000	~$0.279/kWh	Moderate- Fuel dependent	$0.4/kWh
EdL Tariffs	$0.27/kWh				

Conclusion

In this paper, best practices for installation and maintenance of PV and wind power systems were explored, as well as developments in AI and machine learning to improve cleaning and maintenance of these systems and increase their efficiency and output power, and should serve as the foundation for the development and the regulation of this sector in Lebanon.

This paper also looked at Lebanon's potential in the solar and wind energy sectors. Lebanon's persistent electricity problems were aggravated by the country's financial crisis and the total collapse of EdL. Lebanon's target of reaching 30% reliance on renewable energy by 2030 puts the country at a major crossroads for the need for immediate adoption of both solar and wind energy. Political, legislative, and financial accord needs to take place to lay the foundation for public and private initiatives to transpire.

A look at the economic comparison of the cost of power obtained from diesel generators compared with solar and wind power shows the lower LCOE of renewable power sources, and the need to quickly adopt widespread government measures to move towards these renewable sources.

Figure 32 shows the decrease in cost per kW for solar and wind energy between 2010 and 2023, with the LCOE for both solar and wind energy falling dramatically.

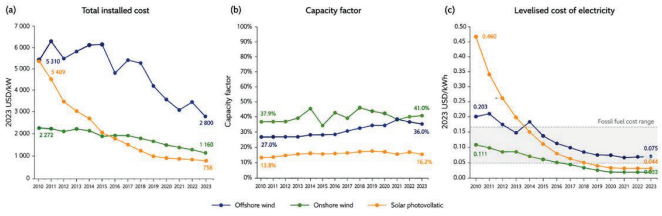


Figure 32: Cost, capacity, and LCOE of renewable energy from 2010 to 2023.[41] Graphs show the change in cost, capacity, and levelized cost of electricity per type of energy between 2010 and 2023.

Lebanon's current economic instability and financial paralysis, worsened by the prevailing security situation, hinder the progress of any sizable transformative renewable energy projects and impede further development of this sector. Regulations have been put in place, but not yet made effective, awaiting the establishment of the Electricity Regulatory Authority (ERA).

The significant growth in recent years in the large number of decentralized PV solar system initiatives in Lebanon, primarily in residential buildings, aims at mitigating the lack of electricity by EdL and to offset the large cost of electricity obtained through private and neighborhood diesel generators, and reveals the lack of application of best practices for installations and maintenance of PV systems, safety concerns, as well as the lack of inclusivity and the inequity in the access to these PV systems. The installation of solar panels on residential buildings in Lebanon during the growth period relied partly on previously used or low-quality components, and/or damaged batteries. These components do not meet safety regulations and were prohibited by Decree 6997. The lack of inspectors in Lebanon to examine the installation of solar panels and the structural integrity of the rooftops led to widespread chaos in the installations of these systems and safety concerns, observed firsthand during rooftop visits in Beirut and surrounding areas.

Lebanon needs to adopt a centralized initiative to fix its ongoing electricity problem. Based on the LCOE of each of the energy sources available for Lebanon to take advantage of, which include EdL, solar power, wind power, and diesel generators, and taking into account the economic, political and financial environment that Lebanon is currently witnessing, the optimal solution for Lebanon is to implement a hybrid system that focuses on addressing the challenges of EdL, adopting wide scale solar energy projects, putting into motion its wind power initiatives, and relying on diesel generators to cover the remaining gap in electricity supply. By doing so, Lebanon can benefit from uninterrupted electricity, decrease the chaos in installations of PV solar systems, lessen dependence on polluting diesel generators, and attract investors.

AI Acknowledgments

During the preparation of this work, the author used Google Gemini to research industry publications and acronyms. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Acknowledgments

The author would like to thank Professor Nazih Moubayed, Dean of the Faculty of Engineering at the Lebanese University, for his invaluable guidance and review throughout the development of this paper.

References

1. VASyR 2022 - Energy.Pdf. https://ialebanon.unhcr.org/vasyr/files/vasyr_chapters/2022/VASyR%202022%20-%20Energy.pdf (accessed 2026-05-09).
2. Solar-Photovoltaic-PV-Booklet-EN.Pdf. https://files.acquia.undp.org/public/migration/lb/Solar-Photovoltaic-PV-booklet-EN.pdf (accessed 2026-05-09).
3. *Global Wind Atlas*. https://globalwindatlas.info (accessed 2026-05-09).
4. PowerPoint Presentation. https://lcec.org.lb/sites/default/files/2025-10/Solar%20PV%20Status%20Report%202023_VF.pdf (accessed 2026-05-09).
5. Abdallah, M.; Zisis, I. Wind-Induced Loads on Photovoltaic (PV) Panels on Low-Sloped Gable Roofs: Panel Size Effects. *J. Wind Eng. Ind. Aerodyn.* **2025**, *263*, 106129. https://doi.org/10.1016/j.jweia.2025.106129.
6. El Hammoumi, A.; Chtita, S.; Motahhir, S.; El Ghzizal, A. Solar PV Energy: From Material to Use, and the Most Commonly Used Techniques to Maximize the Power Output of PV Systems: A Focus on Solar Trackers and Floating Solar Panels. *Energy Rep.* **2022**, *8*, 11992–12010. https://doi.org/10.1016/j.egyr.2022.09.054.

DOI: 10.36838/IJHSR820.1

7. Al Sayegh, H.; Chowdhury, A. G.; Zisis, I.; Elawady, A.; Estephan, J.; Tolera, A. Full-Scale Experimental Investigation of Wind Loading on Ballasted Photovoltaic Arrays Mounted on Flat Roofs. *J. Wind Eng. Ind. Aerodyn.* **2025**, *256*, 105963. https://doi.org/10.1016/j.jweia.2024.105963.
8. Panjawani, R.; Jain, D.; Bhandari, K.; Gaikwad, S.; Deokar, S. Design and Analysis of Solar Structural and Mountings for Solar Panel. *Int. J. Future Gener. Commun. Netw.* **2020**, *13*, 668–679.
9. Sugianto, S. Comparative Analysis of Solar Cell Efficiency between Monocrystalline and Polycrystalline. *INTEK J. Penelit.* **2020**, *7*, 92. https://doi.org/10.31963/intek.v7i2.2625.
10. *(PDF) A Comprehensive Review of Energy Storage Batteries for Solar Energy Systems*. ResearchGate. https://www.researchgate.net/publication/400788517_A_Comprehensive_Review_of_Energy_Storage_Batteries_for_Solar_Energy_Systems (accessed 2026-07-02).
11. *Monocrystalline vs Polycrystalline Solar Panels | American Solar Energy Society*. https://ases.org/monocrystalline-vs-polycrystalline-solar-panels/ (accessed 2026-04-03).
12. Al Garni, H. Z.; Awasthi, A.; Wright, D. Optimal Orientation Angles for Maximizing Energy Yield for Solar PV in Saudi Arabia. *Renew. Energy* **2019**, *133*, 538–550. https://doi.org/10.1016/j.renene.2018.10.048.
13. Nahim, H. M.; Ibrahim, T.; Faraj, J.; Ali, S.; Khaled, M.; Darwiche, M. Optimizing Solar Panel Tilt Angles for Enhanced Energy Production: Case Studies in Lebanon. *Case Stud. Therm. Eng.* **2025**, *73*, 106720. https://doi.org/10.1016/j.csite.2025.106720.
14. *Comprehensive Analysis of Solar Panel Performance and Correlations with Meteorological Parameters - PMC*. https://pmc.ncbi.nlm.nih.gov/articles/PMC10733962/ (accessed 2026-04-03).
15. Abid, M.; Shafiq, W.; Dayem, M.; Khan, H. Step-by-Step Installation of Solar Energy and Its Maintenance. *Int. J. Emerg. Eng. Technol.* **2025**, *4*, 1–5. https://doi.org/10.57041/7pbe1w24.
16. Emeara, M.; Hamdy, E.; Abdel Gawad, A. Hybrid Renewable Energy System for a Sustainable House-Power-Supply ---------- ---------- "Journal of Advanced Research in Fluid Mechanics and Thermal Sciences." **2021**, *87*, 91–107.
17. *Sun, sensors and silicon: How AI is revolutionizing solar farms | World Economic Forum*. https://www.weforum.org/stories/2024/08/how-ai-can-help-revolutionize-solar-power/ (accessed 2026-04-03).
18. Vasiljev, P.; Borodinas, S.; Bareikis, R.; Struckas, A. Ultrasonic System for Solar Panel Cleaning. *Sens. Actuators Phys.* **2013**, *200*, 74–78. https://doi.org/10.1016/j.sna.2013.01.009.
19. *(PDF) Generative AI for Predictive Maintenance in Renewable Energy: Wind Turbines, Solar Farms, and Beyond*. https://www.researchgate.net/publication/395833762_Generative_AI_for_Predictive_Maintenance_in_Renewable_Energy_Wind_Turbines_Solar_Farms_and_Beyond (accessed 2026-04-03).
20. Sepúlveda-Oviedo, E. H. Optimizing PV Maintenance: Methods, Cleaning Frequency, and a Selection Protocol. *Energy Rep.* **2025**, *14*, 1578–1605. https://doi.org/10.1016/j.egyr.2025.07.008.
21. Koondhar, M. A.; Albsha, L.; Mahmoud, S. A.; Amin, L. G.; Mahariq, I. Comparative Analysis of Onshore, Offshore, and Floating Wind Turbines for Renewable Energy Systems: A Review. *Ocean Eng.* **2026**, *343*, 123243. https://doi.org/10.1016/j.oceaneng.2025.123243.
22. Caduff, M.; Huijbregts, M. A. J.; Althaus, H.-J.; Koehler, A.; Hellweg, S. Wind Power Electricity: The Bigger the Turbine, The Greener the Electricity? *Environ. Sci. Technol.* **2012**, *46* (9), 4725–4733. https://doi.org/10.1021/es204108n.
23. *(PDF) Offshore Wind Farms Research Paper*. ResearchGate. https://www.researchgate.net/publication/343878559_Offshore_Wind_Farms_Research_Paper (accessed 2026-04-03).
24. *(PDF) Geotechnical desgin consideration for onshore wind turbine shallow foundations*. ResearchGate. https://www.researchgate.net/publication/326580137_Geotechnical_desgin_consideration_for_onshore_wind_turbine_shallow_foundations (accessed 2026-04-03).
25. *Laying the foundation for wind turbines now and in the future | Wind Systems Magazine*. https://www.windsystemsmag.com/laying-the-foundation-for-wind-turbines-now-and-in-the-future/ (accessed 2026-04-03).
26. Ren, Z.; Verma, A. S.; Li, Y.; Teuwen, J. J. E.; Jiang, Z. Offshore Wind Turbine Operations and Maintenance: A State-of-the-Art Review. *Renew. Sustain. Energy Rev.* **2021**, *144*, 110886. https://doi.org/10.1016/j.rser.2021.110886.
27. Su, B.; Guo, T.; Alam, Md. M. A Review of Wind Energy Harvesting Technology: Civil Engineering Resource, Theory, Optimization, and Application. *Appl. Energy* **2025**, *389*, 125771. https://doi.org/10.1016/j.apenergy.2025.125771.
28. Kaynia, A. M. Seismic Considerations in Design of Offshore Wind Turbines. *Soil Dyn. Earthq. Eng.* **2019**, *124*, 399–407. https://doi.org/10.1016/j.soildyn.2018.04.038.
29. Jiang, Z. Installation of Offshore Wind Turbines: A Technical Review. *Renew. Sustain. Energy Rev.* **2021**, *139*, 110576. https://doi.org/10.1016/j.rser.2020.110576.
30. Beiter, P.; Rand, J. T.; Seel, J.; Lantz, E.; Gilman, P.; Wiser, R. Expert Perspectives on the Wind Plant of the Future. *Wind Energy* **2022**, *25* (8), 1363–1378. https://doi.org/10.1002/we.2735.
31. Kou, L.; Li, Y.; Zhang, F.; Gong, X.; Hu, Y.; Yuan, Q.; Ke, W. Review on Monitoring, Operation and Maintenance of Smart Offshore Wind Farms. *Sensors* **2022**, *22* (8), 2822. https://doi.org/10.3390/s22082822.
32. Julian, M.; Bassil, N.; Dellagi, S. Lebanon's Electricity from Fuel to Solar Energy Production. *Energy Rep.* **2020**, *6*, 420–429. https://doi.org/10.1016/j.egyr.2020.08.061.
33. Dagher, L.; Diab, S.; Zwein, R. From Crisis to Opportunity: Advancing Solar Energy in Lebanon Through Effective Policymaking. *Arab Reform Initiat.* **2025**.
34. *Global Solar Atlas*. https://globalsolaratlas.info/download/lebanon (accessed 2026-04-03).
35. *Details of Decree 6997 on Mandatory Solar Standards | LCEC*. https://lcec.org.lb/node/2990 (accessed 2026-04-03).
36. *Saddi Streamlines Solar Panel Installation Process*. This is Beirut. https://thisisbeirut.com.lb/articles/1311982/saddi-streamlines-solar-panel-installation-process (accessed 2026-04-03).
37. LCECD FIRST DRAFT. https://lcec.org.lb/sites/default/files/2022-08/2020%20Solar%20PV%20Status%20Report%20-%20Final.pdf (accessed 2026-04-03).
38. *YASA Warns of Solar Panel Hazards*. MTV Lebanon. https://www.mtv.com.lb/news/1642554 (accessed 2026-05-16).
39. *ESIA Review Lebanon Wind Power*. NCEA - EN. https://www.eia.nl/en/project-database/esia-review-lebanon-wind-power/ (accessed 2026-04-18).
40. Wind Power in Lebanon the Blades Are Finally Turning.Pdf. https://www.cedro-undp.org/Library/Assets/Gallery/Publications/Wind%20Power%20in%20Lebanon%20the%20blades%20are%20finally%20turning.pdf (accessed 2026-04-03).
41. Renewable Power Generation Costs in 2023.
42. Ahmad, A. *Distributed Power Generation for Lebanon*; World Bank, Washington, DC, 2020. https://doi.org/10.1596/33788.

43. Lebanon Electricity Sector Cost Recovery Plan.Pdf. https://edl.gov.lb/Lebanon%20Electricity%20Sector%20Cost%20Recovery%20Plan.pdf? (accessed 2026-07-05).

Author

Adam Zahran is currently a junior in the International Baccalaureate (IB) Program at International College and aspires to pursue his studies in the field of Engineering, particularly Electrical Engineering.

DOI: 10.36838/IJHSR820.1

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

Exploring The Feasibility of Interstellar Propulsion: Assessing Sub-Century Travel to Alpha Centauri Under Modern Constraints

Advait Y. Reddy

The Aga Khan Academy Hyderabad, Telangana, India; advaitreddyy@gmail.com

ABSTRACT: The pursuit of interstellar travel has long pushed the boundaries of physics and engineering and has tested the limits of what humanity can utilize our resources to engineer a vision to escape the bounds of our solar system to reach nearby stars. The target is Alpha Centauri, which is 4.37 light-years away from Earth. Recent research in interstellar rocket designs, both theoretical and experimental, includes nuclear pulse rockets, nuclear thermal rockets, light sail propulsion, antimatter propulsion, and more designs. This review investigates the feasibility of such interstellar propulsion systems and designs, with practical constraints placed upon them, with an emphasis on the availability of resources and the constraints of physics. Through the amalgam of existing rocket propulsion studies and direct calculations using mathematical equations, this paper evaluates several leading propulsion models against the known limits of material endurance, physics, and energy scalability. Preliminary studies suggest that interstellar travel itself is practical based on physics, but achieving it within a century of travel is a vast aspiration beyond our Earthbound limitations, present-day energy production, and our technological limitations. This paper aims to contribute to interstellar exploration research by identifying pathways for feasible sub-century star transit among existing models and designs.

KEYWORDS: Propulsion Systems, Star Transit, Laser Lightsails, Alpha Centauri, Light Years.

Introduction

Ever since the dawn of our journey to conquer space, humanity has pursued one question above all others: Are we alone? This was a daring question that we ourselves knew when we started our quest, an epiphany predating all our space efforts. In fact, it was in the 1830s that astronomer Thomas Henderson revealed the vast distance between us and the nearest possible source for life, another star,[1] Alpha Centauri, then approximated at 3.57 light-years (ly), though later approximated at 4.37ly, a large distance we might never reach even with the technology of 2026 and beyond. Figure 1 shows how vast these distances really are, using a scaled measure to compare the distance between Earth and Alpha Centauri.

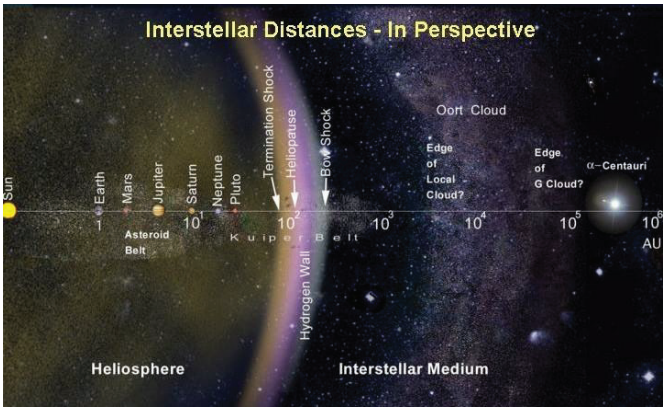


Figure 1: A visualization of the distance from the Sun to A.Centauri. Taken from R.Meraldt & P.Liever, JPL.

But years passed, new ideas came out; some were thought to be absurdities, yet some remain to this day a possibility for interstellar travel, not ruled out by science, formulated by the brilliance humanity has had to offer. These ideas include: nuclear pulse rockets; nuclear thermal rockets; laser lightsails, specifically Breakthrough Starshot;[2] and other concepts[3] which include theoretical/hypothetical models for interstellar exploration, which include antimatter propulsion, warp drives, and many more, which will be explored thoroughly in this paper. Nonetheless, achieving a star transit, especially within our lifetimes, has some requirements. These *requirements* all center on one particular necessity: a speed of around 10% of the speed of light (0.1c), the benchmark for interstellar travel.[4] This would take us to Alpha Centauri in about 42 years or so with minimal time dilation and relativistic effects, a limitation of technology preventing us from traveling at further percentages of the speed of light (>0.3c). As per the principles of relativistic energy proposed by Albert Einstein,[5] it requires immense energy—beyond anything practical, considering the resources required for a small probe to reach even a fraction of the speed of light. Calculated by Einstein's formula for relativistic kinetic energy, determined by $Kinetic\ Energy\ (K.E)=mc^2(\gamma-1)$, where the Lorentz factor (γ) is given by $\gamma = 1/\sqrt{1-(v^2/c^2)}$. To put this into scale, using this formula, for a probe of around 1000 kg or 1 tonne to reach even 10% of the speed of light (threshold for relativistic speeds), it would take the object approximately 4.53×10^{17} joules of energy, implausible at our current state. This paper aims to assess both realistic models and hypothetical models of interstellar vessels, with the final goal of determining

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.2

the feasibility of all these concepts/models, as well as of what resources we have right now and what resources we may have in the future. This will result in the conclusion of whether or not the form of transport assessed can reach Alpha Centauri within a century. Feasibility in this scenario is defined as whether or not a journey between two stars, specifically our Sun and Alpha Centauri, is possible with a balance of theoretical and practical understandings of physics and engineering at play.

This paper demonstrates that laser lightsails are the most feasible interstellar transit method relative to the other models. This is because it has a comparatively lower resource requirement as well as more practical technology involved; both of these reasons will be explained during the course of the paper, specifically in the conclusion.

This paper is going to analyze the methods of "interstellar travel" proposed by previous scientists, aiming to provide a framework for establishing which interstellar travel method or "vessel" would be the best chance humanity has as of right now to venture into interstellar territory. This paper analyzes multiple interstellar transport vessels, including recently developed models such as the Breakthrough Starshot lightsail and Antimatter Propulsion Rockets. It has a detailed and definite analysis of all such models identified, which separates it from most "older" papers and reviews that are not updated with analyses of modern scientific concepts.

■ Comprehensive Secondary Review

The paper utilizes a systematic literature review methodology to critically analyze and assess the feasibility of interstellar travel vessels, both realistic and highly theoretical vessel designs. The paper analyses prior studies that explore the same topic or a topic within the lines of what has been identified. Key papers are sourced (1970s-2026) from databases such as Google Scholar, BIS, AIAA, PDG, and NASA databases. These papers are filtered by their citation metrics, as well as their relevance to the analysis, ensuring that the papers are not only highly relevant but are taken from reputed journals with high credibility, indicated by a high citation count.

The paper also aims to provide analytical improvements to existing studies exploring the feasibility of interstellar travel with modern data calculations taken from other papers. The paper addresses the gaps identified in previous papers by acclaimed researchers such as I. A Crawford,[4] Robert L Forward,[7] and other researchers identified in the paper who are consistently referenced. The paper analyzes a few widespread gaps in identifying the feasibility of interstellar travel, which include:

1) Material Advancements: Considering the improvements in previous analyses with current materials that have increased performance,

2) Modern Developments: The paper keeps in mind modern concepts and ideas newly proposed, such as Breakthrough Starshot (proposed in 2016).

3) Energy Factor: The paper identifies the energy required to produce such vessels and successfully operate them.

4) Hypothetical Models: The paper explores hypothetical vessels, which are rooted in theory rather than practice, such as warp drives, antimatter propulsion systems, and wormhole transit.

The paper aims to use these gaps, paired with identified information in previous papers, to create a new analysis, explaining the feasibility of each possible vessel in relation to one another, and to create a complete analysis.

■ Discussion

This section of the paper analyzes the various propulsion systems and interstellar transit concepts identified as potential prospects for real-life implementation. Each section will use a comprehensive secondary analysis, as well as a detailed overview of whether or not a transit method is practically feasible and mathematically viable.

Laser Lightsails:

This section goes over the concept of laser lightsails, specifically the ones modeled in the project Breakthrough Starshot. The section aims to identify whether this concept is practically feasible, with respect to the availability of resources as of right now and in the future. The mathematical feasibility of laser lightsails will be evaluated in this paper, as well as the resource requirements and theoretical performance for this vessel. A laser lightsail, also known as a laser-pushed lightsail, is an interstellar travel concept that relies on a regular solar sail being accelerated by a gigawatt-scale beam of light, pushing it to nearly 20% of the speed of light (0.2c). A lightsail is composed of 2 major parts: a highly reflective sail about a few nanometers thick, similar in purpose to a sailing ship, and a small probe at the center of the sail where crucial parts like communication gear, navigation, and a miniature camera are placed. The most important part of this design is the sail itself, one that is designed to reflect nearly all the rays of light that fall upon it. It works on the principle of light having momentum (p), as discovered by physicist James Clerk Maxwell and later formulated into an equation by Albert Einstein,[8] given by the formula $E=pc$. The lightsail uses this framework to gain momentum when photons bombard the sail and get reflected, transferring their momentum to the lightsail. However, this process is "too slow" for the lightsail to reach relativistic speeds, with a model lightsail sent into space by the Planetary Society only reaching a maximum of 8.89 km/s or 32000 km/h. This brings in the necessity to increase the number of photons that get reflected by the lightsail, which a laser can do. The light-based laser, composed of a massive amount of photons and hence a massive amount of energy, can accelerate the lightsail to significant speeds using the energy-momentum relation, reaching up to 0.2c, and able to reach Alpha Centauri within 20 years, as shown in Figure 2.

The following sections will identify the various hindrances that prevent the laser lightsail from being feasible as of right now. These hindrances include:

i) Mass scaling;

ii) Laser dynamics and sail material challenges;

iii) Communications, navigation, and camera systems;

iv) Collisions with space dust;
v) Lack of a braking system.

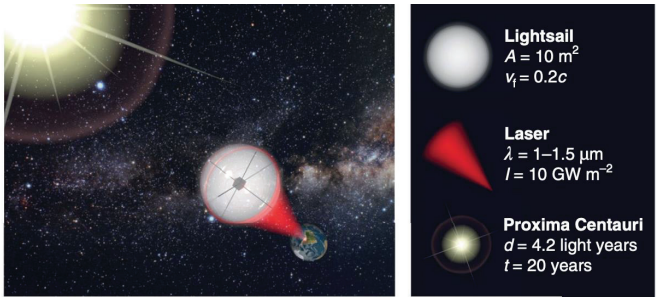


Figure 2: Vision for the Starshot nanocraft. A phased array of lasers, or a light beamer, will propel the nanocraft towards Proxima Centauri. The nanocraft consists of the lightsail, discussed in this Perspective, and a nanochip that will collect data to transmit back to Earth. Taken from Harry. A Atwater's paper "Material challenges for the Starshot Lightsail."

Hindrance: Mass Scaling:

Breakthrough Starshot is a gram-scaled model of a laser-lightsail designed for interstellar travel. The reason for this mass scaling is that it requires significantly less energy to accelerate a probe of a lower mass to relativistic speeds when compared to a probe of a higher mass, as dictated by Einstein's general relativity. This mass is the *first* practical hindrance identifiable when looking at Breakthrough Starshot, with the idea of fitting in all navigation, communication, and camera technologies within a gram scale, unrealistic when thinking of current technology available as of 2026. It is extremely difficult for us to design a probe with a total mass of less than 1 kg, let alone a gram-scale probe, with significant technological advancement required to reduce this mass to the required amount of less than 1 gram.[9]

Hindrance: Laser Dynamics and Sail Materials:

The *second* practical hindrance of the Starshot laser lightsail is the phased array laser itself. Requiring a scale of energy too high, the laser array is something not achievable as of today. In this case, the laser would have to deliver an array with a net intensity of about 10 GW m^{-2} across the sail of an area of 16 m^2 for a 100 - second duration, a currently unachievable feat.[9]

The laser system is likely to be an optically dense phase array of many individual laser elements, such as kW-scale solid-state diode laser-amplifiers.[10,11] The laser system would have to sustain a consistent array without overheating, without consideration of the atmosphere absorbing energy from the laser. The laser system has to be precisely controlled, and nanoscale engineering has to be conducted to precisely control the trajectory of the probe in the exact direction at which the probe could escape the solar system and go towards Alpha Centauri. Atmospheric refraction should be taken into account when the transition of the medium from the atmosphere (1.0003) to the vacuum of space (1.0000) takes place. This is to precisely control the direction at which the probe is pointed, ensuring that Starshot can travel precisely towards Alpha Centauri with a very low margin of error (<1%) for the directional travel.

The fact that the laser array is ground-based vastly amplifies scintillation. Scintillation is where a material emits light after absorbing high-ionizing radiation, which, in the case of Starshot, can be a heavy downside. While space-based arrays counter this issue, they increase the price by countless magnitudes. This scintillation causes random intensity fluctuations and beam wander in the laser, distorting the array's focused spot on the probe's sail, which induces lateral accelerations and could potentially curve its trajectory off course. Thermal and material limits also hinder the effectiveness of the array,[10] while the sail is theorized to be a perfect reflector of light; that is not practically possible.[10] This comes into play when considering the large scale of the laser, approximately 10^{16} J or more, with the sail needing to reflect 99.999% of light to prevent the material from heating or being damaged by the concentrated array.[10]

The leading prospect material for the sail is a dielectric mirror, one of the few man-made materials capable of reaching and even exceeding the required reflectivity of 99.999%. The mirror works by layering thin, transparent materials with alternating high and low refractive indexes, causing light waves to interfere with it *constructively*, resulting in extremely high reflectivity (>99.9%) for specific wavelengths.

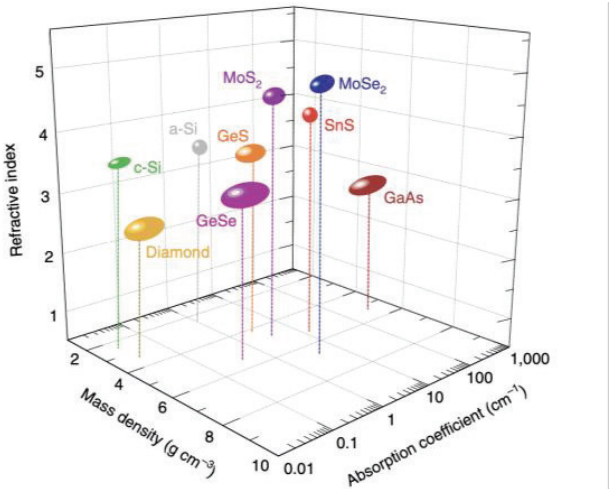


Figure 3: Graph of Refractive Index, Mass-density, and the absorption coefficient of potential materials that can be used for lightsails. Taken from Harry. A Atwater's paper "Material challenges for the Starshot Lightsail."

The drawback with this material is that each layer of the mirror must be a quarter of a wavelength thick, and the mirror itself must have an area of 1 m^2, something that hasn't been achievable to date. Materials research explores new materials like graphene, a material actively being researched by the Breakthrough Starshot team to this day due to its extremely low thickness (0.34-0.35 nm), low mass (0.77 milligram (mg) per m^2), and its high tensile strength, which could potentially withstand the laser array. The problem is that it isn't very reflective at all (0.1-0.3%), meaning that an additional layer of reflective coatings or reflective material would be required for it to be used on a lightsail. Other materials, despite having low absorption coefficients, low mass densities, and high refractive indexes, have tradeoffs, as illustrated by Figure 3 shown above.[9]

Hindrance: Communication, Navigation, and Camera Systems:

The *third* major problem with the concept of a laser lightsail is the communication and navigation system, which, within

DOI: 10.36838/IJHSR820.2

the weight of 1 g, can only communicate weak signals that will be impacted by a Doppler effect due to the high velocity of the spacecraft, causing a shift in the wavelength of the signal emitted.[11] This problem can be solved in 1 major way: making the communication system more powerful (stronger signals), which would normally require a bigger antenna to be used. But this isn't possible due to the strict constraints placed on the mass of the probe, with innovations required in the precision of the signal, a sub-arcsecond precision. Potential solutions include gravitational lensing, which is the use of a star's gravity to focus the signal back towards Earth, laser communication, and a relay concept where a line of probes passes the data between each other, a complex and unlikely alternative. The expected data would consist of images of exoplanets, stars, and scientific measurements (such as radiation, magnetic fields, and gases in space). However, fitting all the required sensors into the camera combined with the navigation and communication systems within a single gram is virtually impossible with present technology. When bringing in the idea of a navigation system, a gyro and a star tracker can likely be placed within a gram's weight, but the idea that they can sustain the high acceleration undergone by the probe, which requires radiation-hardened MEMS that can survive 15000 g, is unlikely.[12] When talking about the camera, while medical equipment for surgical purposes can go beneath the gram-mark, none of them meet the image quality requirements for the Breakthrough Starshot mission, which requires a minimum of 2 megapixels (MP).[13]

Hindrance: Collisions with space dust:

The fourth problem that arises is the likelihood of collision with space dust, most likely to occur with interplanetary dust, a factor only limited by the low exposure time of the probe to this form of dust within the solar system. The problem of collisions lessens significantly when entering the Oort cloud and the interstellar space between our solar system and the Centauri system, with an approximate particle density of 1 atom per cubic centimeter, making collisions unlikely. When setting aside the probability of collisions, an actual collision with interstellar and interplanetary dust could cause surface erosion over the sail itself; the erosion largely varies depending on the dust that it collided against, a factor purely determined by probability we ourselves don't know yet, but one that is for sure not likely. The greatest risk comes with passing through interstellar clouds like the Oort cloud, with an approximate density of 10^9 atoms per km^3, where an impact is not likely but cannot be ruled out.[14]

Hindrance: Lack of a braking system:

The *fifth* and last problem identified in this paper is the lack of braking in the probe,[15] meaning the sail would drift past the Centauri system, leaving only a short window for capturing photos. Despite this issue, the braking system is seemingly impossible to place onto the lightsail, considering the mass requirements, with the use of stars' own radiation/gravity to slow down the probe, or otherwise known as photogravitational assist, theorized to be one of the solutions to this braking problem. The main challenge arising with this is the precision required, especially important here considering the unpredictable forces and radiation at play when encountering the star system, requiring precise calculations, measurements, and orientation to achieve this, something practically improbable considering our lack of knowledge of the Alpha Centauri star system.

Despite these problems, the lightsail has emerged as the premier option for humans to send a fully operating probe to another star system, with most of the problems either already solved or close to being solved, far less "hypothetical" and "impractical" as compared to the other transit options, which will be discussed in this paper.

Nuclear Pulse Rockets:

This section reviews nuclear pulse rockets, determining the current feasibility of the design mathematically and practically based on resource availability. The section will utilize a comparative analysis in comparing the resources required for this project to go interstellar with the resources we currently have right now.

Nuclear Pulse Rockets, as envisioned in Project Orion,[16] propel spacecraft by detonating small nuclear fission/fusion bombs behind a pusher plate to generate thrust from expanding plasma. The rocket works by detonating the bomb, which releases massive amounts of neutrons and X-rays that heat the propellant to millions of degrees in Kelvin.[17] This propellant is instantly vaporized and ionized, releasing plasma that expands at 19-31 km/s in the form of a "fireball," which creates impulse when striking the pusher plate and accelerates the rocket. The "fireball" accelerates at 10000 g peaks, which is theoretically too great for a pusher plate to handle. Due to this, a shock absorption mechanism is implemented in the rocket, which utilizes hundreds of hydraulic absorbers to convert these 10000 g peaks into a constant 1-4 g acceleration, distributing force evenly across the vehicle mass. Figure 4 shows the implementation of these systems into the rocket, with a detailed structure of where these mechanisms are implemented into the rocket, helping understand the thrust process through visual demonstration.

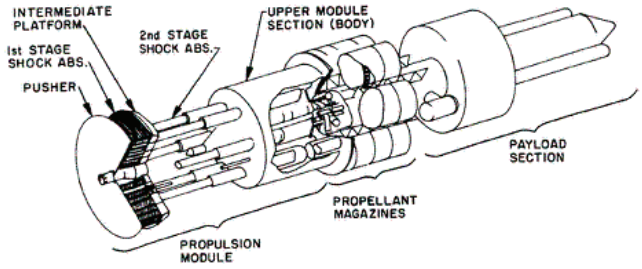


Figure 4: The image above is a detailed structure of a nuclear-pulse rocket with numerous individual components identified, such as the propellant magazines and the payload section, among others. Taken from an article from the "Department of Electrical and Electronic Engineering" titled "nuclear pulse propulsion."

Concepts of nuclear pulse rockets enable interstellar travel by consistent thrust and impulse through controlled nuclear explosions, helping the rocket reach fractions of the speed of light, theoretically up to 10%, but practically up to 5% in ideal conditions.[18] Fission rockets are capped at about 20-100 km/h,

whereas fusion-based rockets range from 0.01-0.1c. This is primarily due to the much higher plasma exhaust velocities, occurring due to fusion's higher energy density and lighter reaction products. Plasma velocity scales as $v_{ex}\sqrt{2E/m}$ where ve is the exhaust velocity of plasma.[19] The exhaust velocity is much higher for fusion reactions because the mass of the ions ejected/yielded is much lower compared to fission products from U-235 or Pu-239. This can be shown by the reactions themselves, with the example of deuterium and tritium fusion with Lithium-6 Deuteride (LiD) as a fuel being compared to U-235 fission as shown below:

Fusion: $^{6}Li + {}^{1}n \rightarrow {}^{4}He + {}^{3}H$ + energy (fission reaction to produce tritium)

$^{2}H + {}^{3}H \rightarrow {}^{4}He + {}^{1}n$ + energy (Deuterium-Tritium Fusion)

Fission: $^{235}U + {}^{1}n \rightarrow {}^{141}Ba + {}^{92}Kr + 3{}^{1}n$ + Energy (splitting of uranium nucleus)

These reactions show that fission reactions have ions of much higher masses (Ba-141, Kr-92) when compared to Alpha particles released by fusion reactions, allowing for the exhaust velocity of fusion reactions to be much higher than that of fission reactions.

This can be further shown by the Tsiolkovsky rocket equation,[19] which calculates the maximum velocity achievable by a propellant-based rocket, given by $\Delta v = v_e \ln(m_0/m_f)$, utilizing mass ratios to show why fusion reactions are preferred, where mo is the total mass of the rocket and mf is the final total mass. When comparing fusion-based and fission-based rockets, fusion has a much higher exhaust velocity (v_e) as previously assessed. Now, comparing the two of them, with the final velocity of fission being 25 km/s and for fusion being 3200 km/s, the mass ratio (m_o/m_f) in the Tsiolkovsky rocket equation required for fission rockets is 10^{120}, and the mass ratio required for fusion rockets is 6.52, the latter of which is the only feasible possibility. Assuming that we would want the rocket to travel at 0.01c, which would take approximately ~440 years for the fusion rocket to reach Alpha Centauri, whereas it is impossible for the fission rocket. The energy required for the rocket to travel at such speeds can be calculated using the formula for Kinetic energy, $E_k = (1/2)mv^2$, where v is the exhaust velocity of plasma. A total propellant mass of 5.2×10^9 kg and a velocity of 0.01c is needed for an interstellar-scale version of the nuclear pulse, with 2.3×10^{21} J of energy required for this. A vessel like this would be "city-sized", as described by Project Orion itself.[16] The amount of energy itself is practically implausible, with a scaled estimate of around ~ 1 million nuclear fusion bombs required to power this gargantuan vessel, each weighing approximately 5 tons. For instance, we only have around 10,000 nuclear fusion missiles on Earth today. Not to mention that the scale of a vessel like this requires immense structural engineering, with the engineering required far beyond our capabilities today. The pusher plate would also have to be able to sustain the plasma "fireballs" expanding at considerably high speeds for nuclear pulse rockets to be a viable form of interstellar travel. The pusher plate aspect is considered impractical as of today due to the lack of any such materials on Earth which can sustain extreme temperature plasma (1,000,000 K+) expanding at rapid velocities.[17] However, if ever this design came to fruition, ignoring all of the impractical aspects of the design, it could be viable with the use of cryogenic sleeper ships to sustain life for the extremely long travel time of 440 years. It is also possible to send unmanned missions using this form of transit. The travel time could also be improved if we could increase the velocity of the ship from 0.01c to 0.1c, reducing it to almost 44 years with the help of relativistic speeds and massive energy scaling. However, another reason this mission isn't likely to happen is the 1963 nuclear test ban,[20] which disallows the use of nuclear bombs in active testing, a major necessity for this program to become experimental, showcasing that political reasons could be at the forefront of why nuclear pulses

This leads to the conclusion that nuclear pulse rockets are currently unviable for *sub-century* transit times to Alpha Centauri, as well as transit to Alpha Centauri on a general basis within the next millennia, when considering the practicalities of the large-scale energy requirements to power the vessel.

Nuclear Thermal Rockets and Nuclear Fusion Rockets:

This section analyses the feasibility of nuclear thermal rockets and nuclear fusion rockets, aiming to identify their practical and mathematical feasibility. It aims to identify the amount of fuel required for these rockets to successfully go interstellar, using that constant to identify whether or not the concept is feasible as of right now.

Nuclear Thermal Rockets:

A Nuclear Thermal Rocket[21] (NTR) is a rocket ship that uses a nuclear reactor to heat a liquid propellant, typically liquid hydrogen, which gets heated up to extreme temperatures and rapidly expands through a nozzle to create highly efficient thrust by expelling the gas at high exhaust velocities (v_e). This is shown in Figure 5, as shown below, showing how plasma is expelled through the nozzle and how the process is undergone. The primary use of the NTR is to provide a more efficient alternative to chemical rockets via the use of nuclear fission for heat instead of chemical combustion. Liquid Hydrogen is among the most efficient fuels available, consistently performing better in terms of specific impulse compared to other rocket fuels such as liquid oxygen and liquid methane.

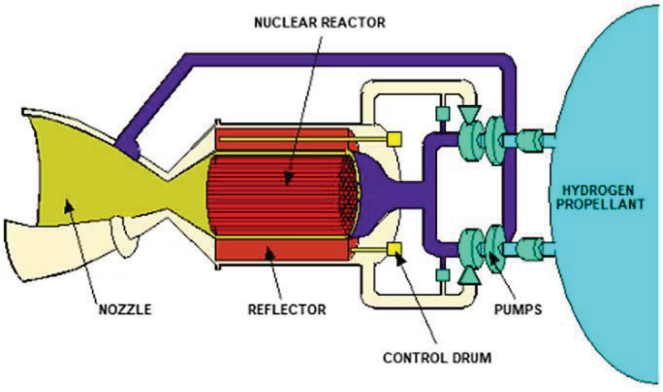


Figure 5: A detailed model of an NTR rocket and its parts aimed at depicting how the rocket generates thrust. Taken from an article by "The Conversation" titled "Nuclear rockets could travel to Mars in half the time – but designing the reactors that would power them isn't easy."

DOI: 10.36838/IJHSR820.2

However, an NTR cannot scale up to being an interstellar vessel because liquid propellant cannot achieve high enough exhaust velocities to have interstellar capabilities. As per the Tsiolkovsky rocket equation, it cannot reach the desired velocities to make interstellar travel possible. This is shown by the calculations below:

$\Delta v = v_e \ln(m_0/m_f)$

Taking the velocity as 9 km/s with a velocity of 0.1c.

$(m_0/m_f) = 10^{120}$

This mass ratio is impossible with current technology and likely even with future technologies. This makes the NTR an unrealistic and implausible vessel for interstellar transit, rather emphasizing its use for interplanetary travel, which is a realistic possibility. Though the energy density of fuels drastically limits the possible expansion of propellant-based rockets to interstellar levels, propellants in the future with GigaJoule/Kg energy densities have the potential to be used for interstellar transit, though none have been discovered as of yet and most likely will not be discovered in the near future. This is also its intended usage, with NASA having plans for using the vessel for Mars transit, among other possibilities.

Nuclear Fusion Rockets (Daedalus Project):

While a traditional propellant-based NTR cannot go interstellar, another nuclear-rocketship was proposed to. That is Project Daedalus,[22] a nuclear fusion rocket using deuterium fusion as the core propellant and using nuclear reactors to fuse them at high temperatures to expel helium gas at high exhaust velocities along with plasma. The rocket contains massive spherical tanks containing highly pressurized helium-3 gas and deuterium gas, as shown in Figure 6.

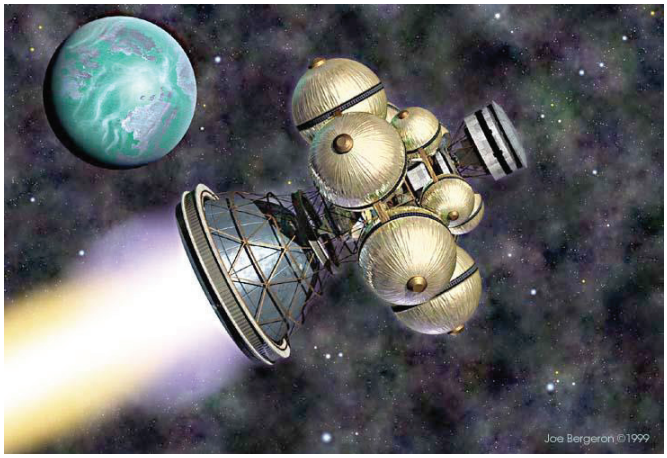


Figure 6: A hypothetical visualization of the Daedalus vessel traveling in space. Taken from an article by "The Velvet Rocket" titled "Project Daedalus."

While a Daedalus rocket is scientifically feasible with viable mass ratios as per the Tsiolkovsky rocket equation, it requires an extremely large amount of deuterium and helium-3 fuel, approximately 50000 tons.[24]

This is practically impossible to gather on Earth due to the scarcity of helium-3, so scientists proposed mining this fuel on Jupiter, a gas giant with abundant sources of helium-3. This is the first feasibility challenge of Project Daedalus, as space mining is a concept of the far future, one that we have not implemented successfully. The scale of space mining is also a massive limitation when considering the requirement of over 50000 tons of fuel. The mining would also take a large amount of time, almost 20 years.

The second challenge is the necessity for pellet microfusion or inertial confinement fusion (ICF),[23] where pellets containing deuterium and helium-3 are compressed and heated to extremely high temperatures using precision lasers to drive fusion.

The lasers would have to be extremely precise for this task, with perfect timing required and high energy recovery, the latter of which has a 17.5% burn-up, or in other words, meaning that only 17.5% of the deuterium and helium-3 undergo fusion.[23] While the remaining 82.5% is unburnt propellant, which still provides momentum to the rocket but doesn't add to the fusion energy of the rocket. In lab conditions, ICF can achieve a maximum of 2-4% burn-up, vastly different from the conditions of a space vacuum. This results in even lower practical burn-up rates, less than 1% when considering the massive scale-up from nanoscale experiments to space vessels operating on kg-scales.[24]

The vessel could reach approximately 12% of the speed of light and was targeting Barnard's star,[22] a target it could reach in 50 years, whereas it could reach Alpha Centauri in about 40 years. All said, the immense scale of this vessel combined with the impracticalities of its fuel usage renders the concept currently impractical; however, with future advancements in helium-3 production, laser internal confinement fusion (ICF), and rocket scaling, we may be able to go interstellar with the Daedalus fusion rocket.

Antimatter Propulsion Systems:

This section aims to analyze the feasibility of antimatter propulsion, identifying the concept's feasibility through the calculation of the mass of antimatter required for rocketships to go interstellar at particular velocities. The review will use this calculation of the antimatter mass required to assess the feasibility of the concept; this is because antimatter can only be obtained in limited quantities, with the quantity being the primary requirement for propulsion systems.

"*Antimatter*" is defined as the mirror image of normal matter,[25] the matter that surrounds us everywhere today. Normal matter consists of protons, electrons, neutrons, quarks, neutrinos, and more.[25] Antimatter particles have the same mass as normal matter particles but have opposite charges and different quantum numbers (baryon and lepton numbers). Our solar system and the entire observable universe seem to be composed overwhelmingly of normal matter, rather than antimatter, a problem or "paradox" that has puzzled physicists for almost a century.[26] This was due to the fact that most physicists believed that when the Big Bang occurred, antimatter and normal matter should've been equally distributed, which isn't the case as per our current identifications. When Antimatter and Normal Matter meet or "collide", they "annihilate" and undergo conversion into pure energy in the form of high-energy photons, neutrinos, or other particles.[27] This can be modeled by Einstein's energy-mass equivalence, $E=mc^2$, where the energy

released by annihilation can be measured by multiplying the combined mass (m) of the colliding particles by the speed of light squared (c^2). This comes into play for propulsion systems when considering the immense energy produced by annihilation compared to its mass. For a gram total of antimatter and matter, $9 \times 10^{\wedge 13}$ joules of energy are produced in annihilation. Comparatively, one gram of enriched nuclear fuel composed of U-235 releases $8.22 \times 10^{\wedge 10}$ joules of energy, and one gram of hydrogen fuel releases $1.2 \times 10^{\wedge 5}$ joules of energy. Due to this significant energy density, matter annihilation is a huge prospect for space propulsion systems, offering more advantages than any other propulsion system available to us.

When talking about propulsion, annihilation becomes a prospect when considering interstellar transit because it is one of the few propulsion systems— if not the only one— that can significantly scale up spacecraft velocities to relativistic scales purely based on meeting energy requirements. The annihilation reaction most viable for space travel is proton-antiproton annihilation, where protons of positive charge collide with antiprotons of negative charge. The reaction is shown below:

$p + \bar{p} \rightarrow 2\pi^- + 2\pi^+ + \pi^0$, where π represents a pion, a primarily unstable particle which then further decays into electrons and positrons, gamma rays (photons), and neutrinos. The primary energy produced by these comes from gamma rays, which are high-energy electromagnetic waves. When used in propulsion, the electrons and positrons produced move at relativistic speeds of about 0.2c to 0.3c; this provides magnetic nozzle thrust in propulsion when these particles are fired out through the nozzle, accelerating the rocket. This can be put into perspective using the relativistic rocket equation,[28] as shown below:

$\Delta v = c tanh^{-1}(v/c)(\ln(m_o/m_f))$

where Δv is the final velocity, and v is the exhaust velocity. Using this formula, the required mass ratios can be identified to produce speeds of 0.1c, 0.25c, and 0.50c. For 0.1c, the mass ratio of m_o/m_f is 1.49, which, in scale, is approximately 24.5 kg of antimatter and 24.5 kg of matter. At this speed, it would approximately take 86 years to travel to Alpha Centauri. For 0.25c, the mass ratio is 2.78, which, in scale, is approximately 89 kg of antimatter and 89kg of matter. At this speed, it would take approximately 34.4 years to travel to Alpha Centauri. For 0.5c, the mass ratio is 9.0, which, in scale, is approximately 400 kg of antimatter and 400 kg of matter. At this speed, it would approximately take 17.2 years to travel to Alpha Centauri.

While Antimatter Propulsion would allow us to reach Alpha Centauri within a century of travel time, it is a technology beyond what we could achieve right now. This is primarily due to the 3 reasons. The first is the lack of production capabilities,[29] with CERN, one of the only facilities in the world capable of producing antimatter, and the largest one, only producing 1-2 nanograms of antiprotons annually, at an efficiency of 10^{-9}. This means that 10^{27} joules of energy would be required to produce just a gram of antimatter, let alone the 24.5 kg required for the bare minimum required for sub-century interstellar transit. The second issue is the lack of storage technology, with penning traps only being able to hold micrograms for minutes at a time via magnetic fields.[30] The environment in space, due to charge effects, can cause these antiprotons to leak at larger densities and scales; with kilogram requirements posing significant storage risks due to annihilation from antimatter leakage. The third and last issue is the lack of thrust efficiency, with 40% of energy escaping as unstable gamma rays,[31] combined with pion beams, as well as muons escaping the nozzle before decay, reducing practical efficiency to about 50%. This results in an increase in mass ratios to accommodate this inefficiency, leading to almost double the amount of antimatter required to produce theoretical thrust, alongside the increased scaling requirements and containment requirements, making it highly unlikely that annihilation thrust can be efficient. This is further reinforced by Figure 7, as shown below, which shows how likely protons and antiprotons are to annihilate at different collision energies. It matches the data shown above by showing that low energies or slow particles give high annihilation rates with pion beams for thrust, as shown by the 0.2-0.3c speed of relativistic electrons and positrons. High energies make annihilation rates low, explaining why CERN needed vast power to produce a tiny yield of antiparticles. This can be explained by quantum field theory (QFT), which states that protons and antiprotons are bundles of quark fields.[32] When quarks and antiquarks touch, they cancel out via the strong force of gluons, releasing energy that forms pions. This explains why cancellation is easier at low speeds and harder at higher speeds, explaining efficient thrust from low annihilation or low speeds, and also why accelerators struggle to make antiprotons.

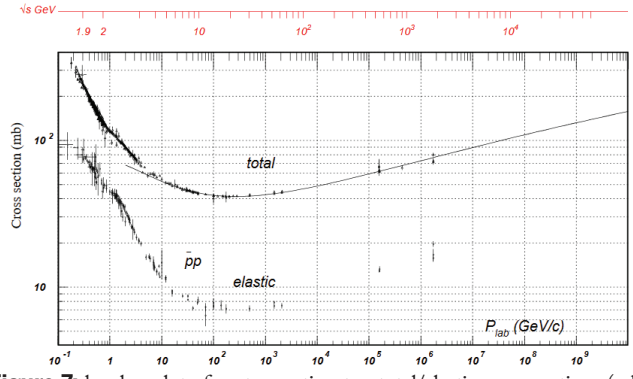


Figure 7: log-log plot of proton-antiproton total/elastic cross-sections (mb) vs energy squared (GeV2). Taken from an article by "The Physics Stack Exchange" about quantum field theory in antiproton-proton collisions.

Solutions for the problem of producing antimatter are often difficult to find, but one proposed solution is to harvest antimatter from space,[33] which is seemingly abundant in antiprotons formed from cosmic rays and solar flares, found in trace amounts and low densities. Scientists have theorized finding large amounts of antiprotons in Jupiter, while unlikely due to the fact that antimatter annihilates very quickly when in contact with normal matter. This solution would require us to find antimatter hotspots in deep space, with abundant spots in the center of our galaxy, as well as spots near supernovas. This, alongside development in our particle accelerators, especially in improving the efficiency in producing antiprotons and antiparticles, is key to increasing our antimatter production to reach a threshold where we can produce enough for interstellar transit, among other potential uses.

DOI: 10.36838/IJHSR820.2

Faster Than Light (FTL) Travel:

This section of the paper assesses the feasibility of hypothetical frameworks for interstellar travel, including wormhole transit and Alcubierre warp drives. The review will analyze how "theoretical" these concepts are, with a practical assessment of whether or not these concepts can become viable in practice, rather than just being rooted in scientific theory.

Faster Than Light (FTL) travel[34] is a theoretical and hypothetical framework in which spacecraft can travel at speeds greater than the speed of light (c) of 3 x 10^8 m/s in a vacuum. FTL is currently considered impossible via relativity, where objects would need to have imaginary mass and infinite energy. This violates causality, which states that an effect cannot happen before its cause, with the violation occurring because FTL travel can allow time travel into the past, causing paradoxes such as the grandfather paradox.[35] The grandfather paradox is a thought experiment in which a local contradiction is highlighted: if you travel back in time to before your father was conceived by your grandfather and kill your grandfather, you would never have been born, meaning you couldn't have traveled back in time to kill your grandfather in the first place, creating an impossible-to-imagine loop or "paradox". FTL travel would allow us to reach Alpha Centauri and beyond, in near-instantaneous times, due to the nature of FTL travel exceeding the speed of light, allowing for infinite speeds.

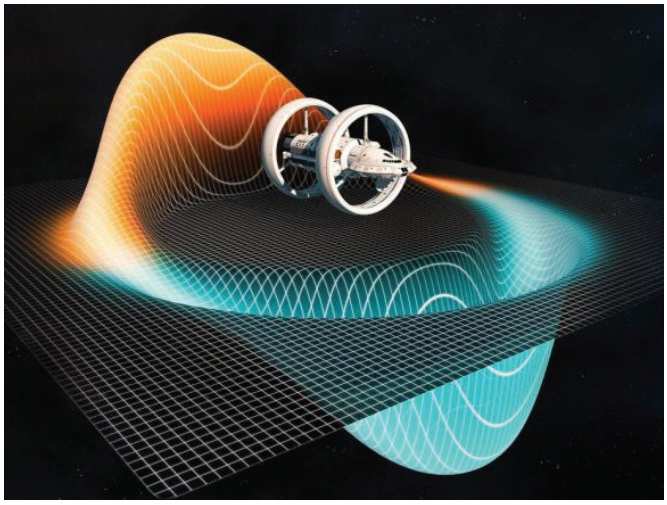


Figure 8: A visualization of the Alcubierre warp drive, contracting Space at its front and expanding it at the back, "warping space-time". Taken from an article about the warp drive by the Space Newsletter.

The main forms of FTL travel proposed by researchers include the Alcubierre Warp Drive[36] and Wormhole Travel.[37] The Alcubierre drive proposes that space-time can be warped, contracting space in front of the ship and expanding it at the back of the ship, creating a "warp bubble" that carries the ship along FTL speeds as shown by Figure 8. While mathematically consistent with Einstein's general relativity, it requires exotic matter[37] with negative energy density and immense energy, negative in nature.

The warp drive is often conceived in science fiction, without much practical experimentability, due to its nature of violating the fundamental laws of relativity and causality.

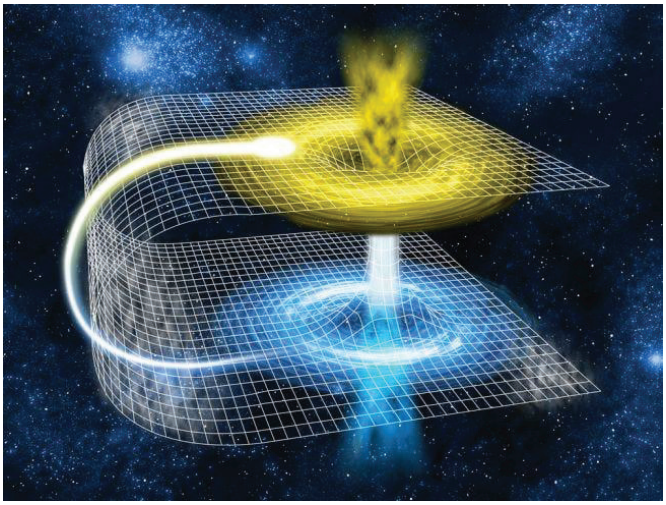


Figure 9: A visualization of a wormhole creating a hole in space-time, allowing for a means to travel between two points in 3D space directly. Taken from an article by "Medium" about the latter.

[22]The other solution is wormholes, which act as shortcuts through space-time, essentially creating a direct path between two discrete points in space-time without traveling to them in a direct path, as shown by Figure 9. General Relativity suggests that they would collapse instantly and that they would need exotic matter with infinite energy to be stabilized.[36] Wormholes are also not proven to exist, with no observational proof, further adding to their "hypothetical" state of conceivability. Wormholes are often present in science fiction rather than real-life experiments and studies, due to their requirement of exotic matter, another term unproven.

Summary of discussion and comparison between interstellar vessels:

This section of the paper is the summary of all the analyses of the "feasibility of different interstellar space travel vessels and the methods they utilize". This paper hypothesizes that laser lightsails are the most feasible interstellar transit method relative to both realistic and theoretical models due to the resource scaling and technological requirements being much lower relative to these systems.

Table 1: A comparison of the energy requirements of different interstellar vessels and their respective transit times. Made with the help of the findings in this paper.

Vessel	Energy Requirement/J	Transit Time
Laser Lightsail (starshot)	10^11	20-50 years
Nuclear Pulse Rocket	10^16	40-80 years
Nuclear Fusion Rocket (daedalus)	10^20	50-100 years
Antimatter	10^17	100-150 years
Warp Drives	-10^44 (negative energy)	Instant-Weeks
Wormholes	Infinite	

The hypothesis restated above holds after the analysis conducted. This is because the laser lightsail is not only the vessel that needs the least energy to operate when compared to other models, as shown above by Figure 10; it is also the most practically applicable model among the lot. This is because much of the laser lightsail's design is grounded in realistic modeling, without any design being explicitly implausible, with the aspect

of the laser array being a bit unrealistic due to the significant energy requirements, but something that can be achieved in the future with significant advancements in energy production. With the development of the laser, the development of reflective sail technologies is advancing, already being developed in the form of a dielectric mirror; this will likely be achieved in the near future with significant scientific advances in the study of reflective materials. The other unrealistic aspect is the controlled mass of only 1 g, with the lack of sufficient parts to build a navigation system, communication system, and a camera being evident. With these implementations in order, the lightsail is proven to be capable of achieving interstellar travel to Alpha Centauri within 20 years at a velocity of 0.2c, grounded in physics and accurately modeled far less than the sub-century transit requirement posited as a metric of "feasibility" in the introduction. Even though it isn't likely that we achieve a fully operating model of Starshot within our lifetimes, it isn't something too unrealistic. Especially when compared to the other interstellar vessel designs, which is the next part of the discussion. While significant advances are required in the field of science to create a fully functional and operating laser lightsail as described in Breakthrough Starshot, these advances are significantly more rooted in realism when compared to nuclear pulse rockets, nuclear fusion rockets, nuclear thermal rockets, and antimatter propulsion systems.

To start, nuclear pulse rockets are completely unrealistic due to their substantial size and scale, paired with the requirement of up to a million nuclear fusion bombs. This count is not only hyper-unrealistic in terms of our current production capabilities, but also unrealistic in terms of the amount of uranium and plutonium deposits present on Earth, not enough to build a million nuclear fusion bombs. This also pairs with our lack of infrastructure capabilities to build a "city-scaled" vessel, which is required to pursue interstellar destinations such as Alpha Centauri. Alongside this, we would have to discover new materials capable of withstanding high-temperature plasma constantly barraging against them to ensure a working pusher plate is designed to allow the rocket to work. All these factors combine and make a concise point against the possibility of nuclear pulse rockets being used for interstellar travel, due to the implausibility in its design, even though they can theoretically reach Alpha Centauri within a century.

Secondly, Nuclear Thermal Rockets (NTR) are chemical rockets using a fusion or fission reactor to generate thrust by rapidly heating liquid hydrogen. The fact that it is a chemical rocket itself proves that the design isn't feasible for interstellar travel, especially interstellar travel within a century's duration. This is because the liquid propellant cannot produce enough exhaust velocity to accelerate the rocket to the required relativistic speeds. While not shown in the figure above, no NTR rocket will ever be capable of achieving interstellar travel due to the combustion-oriented process of flight.

The other nuclear rocket considered is the nuclear fusion rocket, as imagined in Project Daedalus. The reason this design is also impractical and not feasible is that it relies on laser inertial confinement fusion (ICF), a process that has less than 1% efficiency in terms of fusing deuterium and helium-3 molecules into sufficient gas to create enough thrust, rendering the exhaust velocities too low to even consider going interstellar. The other reason being that there is no possible way to generate enough helium-3 on Earth to power this rocket, with the idea of space-mining/harvesting being considered but implausible due to insignificant advancements in the field of space harvesting in the current day.

When talking about antimatter propulsion, it has significant groundwork and theory laid upon it to prove that it is possible, with antiproton annihilation creating more than enough energy to take us to Alpha Centauri if enough antimatter is accumulated. Only 89 kg of antimatter is required to take a spaceship to 0.25c. However, when you look at the key term only, it reveals a lot, because we are only able to produce a few nanograms of antimatter on an annual basis. While antimatter propulsion is the most energy-dense form of propulsion available to us, we simply can't get our hands on enough antimatter to make this system possible. Even if we can scale up antimatter production with the increase in the number of particle accelerators, we wouldn't be able to contain or store the antimatter without risking leakage, which can lead to annihilation and hence major safety concerns. This, paired with the inefficiency of antimatter production, currently at 10^{-9}, makes antimatter propulsion improbable to be implemented in the near future, rendering it unfeasible even though it theoretically can achieve sub-century star transit.

The last option explored, FTL travel — that of the Alcubierre warp drive and Wormhole travel — is heavily rooted in theory but completely neglected in practice. While wormholes and the Alcubierre warp drive do offer us a direct path to a potential interstellar system, they aren't proven, and they might not be proven in the near future, rendering them unfeasible as of right now without any significant scientific advancements in the field. These transit methods require the discovery of exotic matter, matter with negative energies and immense energy, often infinite. FTL travel itself is highly paradoxical in nature, with the specifications of going faster than light with an object of mass impossible as per general relativity, allowing an object to go back in time, creating paradoxes like the grandfather paradox, violating causality, and hence unfeasible as of right now.

This analysis shows that while all of these designs may be unfeasible as of right now, the laser lightsail is the most practical because it is the least impractical, which proves the hypothesis stated in this paper's introduction to be correct, based on the extensive analysis conducted on each interstellar transit vessel.

■ Conclusion

Continuing directly from the discussion, the laser-lightsail has been identified as the most effective and "practically feasible" interstellar travel vessel. Restated from the discussion, the laser-lightsail is the most "feasible" because it has the least energy requirements relative to the other vessels (referenced in Table 1), alongside having the most realistic material requirements, while also having the quickest travel times when compared to the other vessels.

DOI: 10.36838/IJHSR820.2

However, other key findings include: Antimatter Propulsion could be the future of interstellar travel if we can find a way to produce enough antimatter and store it. The mass requirements are around 24.5 kg to reach Alpha Centauri in a sub-century time of 89 years. Nuclear Pulse Rockets remain largely impractical due to the large quantity of nuclear-fusion bombs required, on the scale of millions. This is further perpetuated by the fact that no pusher plate can withstand the extreme temperatures of the plasma produced (upwards of a million K) from the fusion process undergone in the rocket; While contained fusion rockets expelling fusion byproducts as propellant, such as Project Daedalus, are mathematically consistent, they require tons of helium-3 fuel. This is not achievable, as we cannot produce large amounts of helium-3 fuel, meaning that we would have to invent new methods of production. Furthermore, laser internal confined fusion (ICF) is a new–modern nuclear fusion method using highly precise methods that is used in this rocket, however it only has an efficiency of around 2-4%, meaning that we would have to improve the efficiency of laser ICF for this to be applicable; Nuclear thermal rockets (NTR) using liquid hydrogen propellant and nuclear fusion to expel this propellant are impossible for interstellar travel. This is because unobtainable amounts of propellant would be required, exceeding 10^{100} kg, for this rocket to be viable. However, NTR can be feasible for interplanetary travel to nearby planets like Mars, Jupiter, and further beyond, as nuclear fusion is extremely efficient with liquid-hydrogen combustion when covering smaller distances within our solar system. Moreover, wormholes and Alcubierre warp drives are purely theoretical prospects for interstellar travel. While they do not exist and cannot exist according to the laws of physics, as they require exotic matter, if they did, they would certainly be the most efficient process for interstellar travel due to near-instantaneous transit times.

All in all, using this explanation, it is understood that laser lightsails (Starshot) are the most feasible as of today, with Nuclear Pulse Rockets following (fusion-based), then the Antimatter Propulsion System, followed by the Daedalus Project model, after which the theoretical models of wormholes and warp drives enter the frame. This clear order showcases that Laser lightsails, while still not being practical as of today, have the most potential to become a realistic solution in the future, with the rest being less practical comparatively.

To conclude, while interstellar travel remains largely impractical as of 2026, it could be feasible with future advancements in physics and propulsion, which scientists predict could happen by the year 2113, keeping our hopes of going interstellar alive, helping us possibly answer the famed question proposed eons ago: are we alone? All in all, from the extensive analysis conducted, it is highly recommended that researchers look into advancing our research on laser lightsails and the Starshot model, working on improving current material and technology constraints, followed by energy practicalities.

AI Acknowledgments

During the preparation of this work, I used ChatGPT and Claude to find sources for my research and to double-check my calculations. After using these tools, I reviewed and edited the content as needed and take full responsibility for the content of the publication.

Acknowledgments

I would like to thank Professor Andrea Giusti from Imperial College and Mr.Marcin Kedziera from the University of Edinburgh for guiding me throughout my research process and sharing their knowledge in the field of astrophysics and space propulsion. Without their guidance, this paper wouldn't have been possible.

References

1. Falkner, D.E. (2021). Thomas Henderson: Alpha Centauri. In: Stories of Astronomers and Their Stars. The Patrick Moore Practical Astronomy Series. Springer, Cham. https://doi.org/10.1007/978-3-030-80309-4_9
2. Breakthrough Starshot. "Breakthrough Initiatives." Breakthroughinitiatives.org, 2019, breakthroughinitiatives.org/initiative/3.
3. Kaku, Michio. "The Physics of Interstellar Travel." *November* 29 (2002): 166-167.
4. Crawford, I. A. "Interstellar travel: a review for astronomers." *Quarterly Journal of the Royal Astronomical Society, Vol. 31, p. 377-400 31* (1990): 377-400.
5. Einstein, Albert. "On the Electrodynamics of Moving Bodies." Annalen der Physik (1905): 891–921
6. Forward, Robert L. "Feasibility of interstellar travel: A review." *Acta Astronautica* 14 (1986): 243-252.
7. Forward, Robert L. "A program for interstellar exploration." *British Interplanetary Society, Journal* 29 (1976): 611-632.
8. Einstein, A. (1905). Does the Inertia of a Body Depend on Its Energy Content? Annalen der Physik, 18, 639.
9. https://doi.org/10.1002/andp.19053231314
10. Atwater, Harry A., Artur R. Davoyan, Ognjen Ilic, *et al.* "Materials challenges for the Starshot lightsail." *Nature Materials* 17.10 (2018): 861-867.
11. "Breakthrough Initiatives." *Breakthroughinitiatives.org*, 2017, breakthroughinitiatives.org/forum/5?page=6. Accessed 3 Feb. 2026.
12. Tischer, F. "Doppler phenomena in space communications." *IRE Transactions on Communications Systems* 7.1 (2003): 25-30.
13. Parkin, Kevin LG. "The Breakthrough Starshot system model." *Acta Astronautica* 152 (2018): 370-384.
14. "Breakthrough Initiatives." *Breakthroughinitiatives.org*, 2018, breakthroughinitiatives.org/forum/18?page=3. Accessed 3 Feb. 2026.
15. Stern, S. Alan. "Collisions in the Oort cloud." *Icarus* 73.3 (1988): 499-507.
16. Fedi, Marco. "Significant deceleration and mission failure, due to dilatant vacuum, expected for the Breakthrough Starshot program." (2019).
17. NASA. "Orion Spacecraft - NASA." *NASA*, www.nasa.gov/humans-in-space/orion-spacecraft/.
18. Schmidt, G. R., J. A. Bonometti, and C. A. Irvine. "Project Orion and future prospects for nuclear pulse propulsion." J*ournal of Propulsion and Power* 18.3 (2002): 497-504.
19. "Project Orion: A Nuclear Bomb and Rocket – All in One." *Www.centauri-Dreams.org*, www.centauri-dreams.org/2016/09/16/project-orion-a-nuclear-bomb-and-rocket-all-in-one/.
20. Tsiolkovsky, Konstantin E. "The exploration of cosmic space by means of reaction devices." *Scientific Review* 5 (1903): 60-132.

21. CTBTO. "The Comprehensive Nuclear-Test-Ban Treaty (CTBT) | CTBTO." *Www.ctbto.org*, 2023, www.ctbto.org/our-mission/the-treaty.
22. Borowski, Stanley K., David R. McCurdy, and Thomas W. Packard. *Nuclear thermal rocket (NTR) propulsion: A proven game-changing technology for future human exploration missions*. No. E-18199. 2012
23. Yee, Nick. "The Daedalus project." *Retrieved July* 17 (2006): 2006.
24. Betti, R., and O. A. Hurricane. "Inertial-confinement fusion with lasers." *Nature Physics* 12.5 (2016): 435-448.
25. Long, Kelvin F. "Interstellar propulsion using laser-driven inertial confinement fusion physics." Universe 8.8 (2022): 421.
26. Canetti, Laurent, Marco Drewes, and Mikhail Shaposhnikov. "Matter and Antimatter in the Universe." *New Journal of Physics* 14.9 (2012): 095012.
27. Dine, Michael, and Alexander Kusenko. "Origin of the matter-antimatter asymmetry." *Reviews of Modern Physics* 76.1 (2003): 1.
28. Alfven, Hannes, and Oskar Klein. "Matter-antimatter annihilation and cosmology." *Arkiv Fysik* 23 (1962).
29. Ackeret, J. "Zur Theorie der Raketen." *Helvetica Physica Acta*, Vol. 19, 1946, pp. 103–112
30. Cleymans, J., I. Kraus, H. Oeschler, *et al.* "Antimatter production in proton-proton and heavy-ion collisions at ultrarelativistic energies." *Physical Review C* 84.5 (2011): 054916.
31. "Storing Antihydrogen | CERN." *Home. cern*, home.cern/science/physics/antimatter/storing-antihydrogen.
32. Forward, Robert L. "*Antimatter propulsion*." J. Br. Interplanet. Soc.;(United Kingdom) 35 (1982).
33. Peskin, Michael E. *An Introduction to Quantum Field Theory*. CRC Press, 2018.
34. Jackson, G. "Antimatter Harvesting in Space." *NIAC Phase I Final Report, Hbar Technologies, LLC* (2006).
35. Krasnikov, Serguei. *Back-in-time and Faster-than-light Travel in General Relativity*. Vol. 193. Berlin/Heidelberg, Germany: Springer, 2018.
36. Hauser, Jacob, and Barak Shoshany. "Time travel paradoxes and multiple histories." *arXiv preprint arXiv:1911.11590* (2019)
37. McMonigal, Brendan, Geraint F. Lewis, and Philip O'Byrne. "Alcubierre warp drive: On the matter of matter." *Physical Review D—Particles, Fields, Gravitation, and Cosmology* 85.6 (2012): 064024.
38. Morris, Michael S., and Kip S. Thorne. "Wormholes in spacetime and their use for interstellar travel: A tool for teaching general relativity." *American Journal of Physics* 56.5 (1988): 395-412.

Author

Advait Reddy. Y is a young physicist in the works at the Aga Khan Academy Hyderabad in India. He is currently doing the International Baccalaureate (IB) program, seeking to attend a top university abroad to further pursue his aspirations in the field of physics and engineering.

DOI: 10.36838/IJHSR820.2

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

Impacts of Early Intervention Plans on Cognitive and Socio-emotional Development in Down Syndrome Patients

Akanksha Mohanty

Oberoi International School, JVLR, Mumbai, Maharashtra, 400093, India; akm090827@gmail.com

ABSTRACT: This literature review explores the effects of early intervention and its impacts on cognitive and socio-emotional development in infancy and early childhood with Down Syndrome. Being a neurobiological disorder, it affects key areas of cognitive and socio-emotional functioning. As a disorder characterized by its congenital nature, key developmental phases wherein emotional, social, and cognitive skills are meant to mature are inevitably affected. This leads to irregularities in behavioral traits, emotional perception, interpersonal and situational interaction and awareness, and more, which present significant barriers to maintaining a successful professional lifestyle and an emotionally fulfilling personal and social lifestyle. While the structural differentiations are incurable, early intervention plans targeted at the cognitive and socio-emotional consequences of these structural differentiations, implemented in the early stages of Down Syndrome, at infancy and childhood, can support and enhance development, showcasing short-term and long-term benefits. This literature review posits that while there are varied results to early intervention plans and treatment plans that focus on supporting development at early stages to take advantage of the neuroplasticity of the brain at young ages, they are largely beneficial and pose no harm in their contribution to social and emotional development. This literature review expands the in-depth understanding of how and why early intervention plans, in particular, benefit cognitive and socio-emotional development, contributing to the knowledge base of early intervention research for people with Down Syndrome during the infant and childhood stages.

KEYWORDS: Behavioral and Social Sciences, Neuroscience, Down Syndrome, Cognitive and Socio-Emotional Development, Early intervention.

Introduction

Down Syndrome is characterized by challenges with respect to reciprocating someone's feelings, communicating what they think verbally and through body language, and more. Defined as an intellectual disability, it impacts cognitive, socio-emotional, behavioral, and physical development and traits. Phenotypic traits include distinctively flattened facial features and upward-slanting eyes,[1] while behavioral, socio-emotional, and cognitive differentiations include difficulties with maintaining attention, short- and long - term memory, social navigation skills, and executive function.[2] Down Syndrome is a developmental disorder caused by a genetic mutation - Trisomy 21, a biological phenomenon entailing a supernumerary chromosome 21. This occurs when an extra chromosome appears in the 21st autosome, which leads to certain differences in cognitive and socio-emotional development.[3] This mutation leads to certain structural differences in the brain; neurons experience stunted dendritic growth, resulting in a smaller overall brain size, which is considered to play a significant role in cognitive and socio-emotional developmental difficulties. While the biological root cause of these difficulties cannot directly be addressed or cured in any way with current technology and scientific knowledge, early intervention plans target working directly on the hampered skills - for eg. social interaction - and have been known to have beneficial short and long - term benefits, of which improved intellectual and adaptive function (long - term) and higher IQ levels, though these tended to recede slightly with age (short - term) are some of the most significant.[4] These results are primarily the result of early intervention plans, as implementing programs alongside key developmental phases allows this development to be supported and guided at its most crucial and delicate stages. Research into the impacts of responsive teaching and model preschool programs displays considerable improvements in patients with Down Syndrome with respect to cognitive and socio-emotional developmental areas while addressing minor ethical risks, overall outlining the largely beneficial nature of early intervention plans and their valuable contribution to cognitive and socio-emotional development. This literature review analyses the structural differentiations in the brain and provides a baseline association with cognitive and socio-emotional development in Down Syndrome, highlighting the root cause of the underdevelopment of these traits. It then covers cognitive and socio-emotional development in Down Syndrome as a separate topic, establishes the connection between the two more strongly, and introduces responsive teaching and model preschool programs as successful early intervention plans, highlighting their advantages, disadvantages, and risks, and provides an overall view and details the reasons behind the outcomes of early intervention plans. It then evaluates the real-life implications of early intervention programs and how they longitudinally influence the well-being of Down Syndrome patients.

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.3

Discussion

Cognitive and socio-emotional development in patients with Down syndrome:

Cognitive and socio-emotional processes are crucial to effective thought and emotion control. Processes such as executive function (the lack, inhibition, or delay of executive function referred to as executive dysfunction) are found to undergo delayed or insufficient development in patients with Down Syndrome, leading to difficulties in social cognition skills and interpersonal communicative skills. Social understanding and engagement skills are central to intellectual development and a significant driver of overall development,[5] especially during younger years at peak neuroplasticity levels when the brain is still forming and developing cognitive and socio-emotional structures and processes.[6] Many children with IDs (intellectual disabilities) display a difference in rate of development of normal cognitive and socio-emotional abilities rather than a big difference in the development of the abilities themselves,[7] and behaviors exhibiting typical social behaviors occurred at a lower frequency as compared to typically developing children - especially in self-reliant and competitive contexts.[8]

In Down Syndrome, on the other hand, there are multiple early indicators of emerging cognitive understanding that develop atypically, including dyadic interactions between the adult and the infant (mutual gaze, focal nature of gaze, reading of social signals, etc).[5] Early dyadic interactions between infants and adults serve as a critical foundation for socio-cognitive development, interactions which are normally characterized by the infant's steadily growing receptiveness to subtle social and communicative signals.[5] These interpersonal exchanges facilitate the transition to more complex interpersonal interactions within the first year of life, including the emergence of joint attention, intentional pointing, requesting behaviors, and imitative learning, wherein infants assume an increasingly leading role. This is instrumental in cultivating higher-order cognitive skills, most notably emotional recognition and theory of mind,[5] which are essential for successful social navigation and interpersonal understanding,[9] which are also, in turn, big factors of socio-emotional development through further years of childhood. Mutual gaze, an important marker of typical Early dyadic interaction, is slow to emerge in infants with Down Syndrome, but is then maintained at a high level in the middle of the first year,[5] which is when it declines in typical development as the infant begins to pay more attention to the wider social and physical world, perhaps due to a need for more looking time to obtain the same amount of information or hypotonia. The 'focal nature' of the gaze of children with Down Syndrome entails a fixation on the parent's/caregiver's eyes rather than exploring other facial features, which could signify potential future implications in emotional recognition, leading to a significant hurdle in their ability to read critical social cues.[5] The lack of the utmost foundational interpersonal skills compromises intellectual as well as overall development in the brain, as the knowledge garnered from social exploration is not gathered, and development in children with Down Syndrome builds onto already existing manifestations of executive dysfunction. This inhibition of development is mirrored in structural differentiations visible in the anatomical development of the brain, affecting cognitive and socio-emotional development.

The structural differentiations in the brain associated with cognitive and socio-emotional development due to Down Syndrome:

The main basic structural difference exhibited in brains affected by Down Syndrome includes dendritic atrophy, which strongly affects cognitive and socio-emotional development.[10] There are some set criteria against which the extent of dendritic atrophy is measured: the number of dendritic branches, the number of dendritic branch segments, and dendritic length and growth, the underdevelopment of which leads to cognitive impairment.[11]

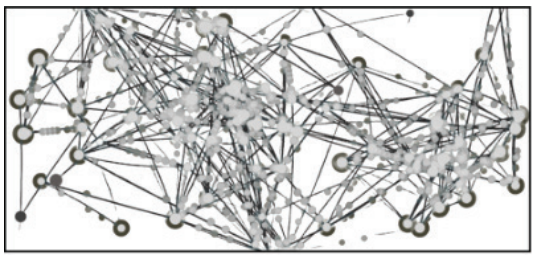


Figure 1: Representation of regular development of neural networks in neurotypical brains showing an approximation of dendritic branching and density.

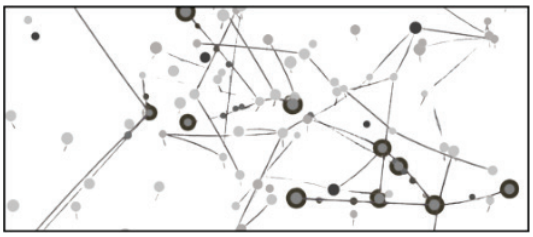


Figure 2: Representation of the development of neural networks in the brains with Down Syndrome, showing an approximation of dendritic branching and density as a result of dendritic atrophy.

The above diagrammatic representations, Figure 1 and Figure 2, demonstrate how brains impacted by Down Syndrome undergo dendritic atrophy, where the number of dendritic branches forming from the neurons reduces, and the frequency of growth decreases, leading to lower dendritic density and weaker neural connections. The neural networks formed in brains that are regularly developing demonstrate significantly higher levels of complexity, interconnectedness, and strength.

A study on the dendritic growth of brains impacted by Down Syndrome was conducted with a group of Down Syndrome patients and a control group ranging from infancy to early childhood stages.[10] While the total number of dendritic branches was higher in infants affected by Down Syndrome, it quickly reduced in comparison to the control group. By around 5 years old, DS patients had a significantly lower total number of dendritic branches as compared to 5-year-olds in the control group, wherein the number had slightly increased with age. A similar effect was seen considering the number of dendritic branches - while infant Down syndrome patients displayed a high number of dendritic branch segments, they quickly fell with age, while unaffected children in the control

group experienced a quick increase in the number of dendritic branch segments.

Dendritic growth exhibited, again, similar tendencies - while there was a visible spurt in growth in the first four months of life, dendritic development significantly deteriorated after. In the control group, dendritic growth and expansion continued through early developmental stages and through childhood, the end age being unclear, but definitely continuing up to the age of five. Dendritic length, however, decreases from infancy itself in direct contrast to the control group, where dendritic length continues to increase as age increases. There is also often a lower number of neurons overall in Down Syndrome patients, the existing neurons solely exhibiting underdeveloped dendritic arborization - a characteristic of old age, earning the label of premature neuron aging. Abnormal dendritic length of the spines and underdeveloped dendritic arborescence are particularly identified in cortical pyramidal neurons, crucial to cognitive functioning.[11] This significantly hinders neuron connection, and the consequent reduction in synaptic density significantly hinders neuron connection, and negatively affects the brain's ability to form proper neural networks. These structural changes in the brain circuitry affect the processing and transmission of information and brain plasticity, including the ability to form new connections and the capacity for learning and adaptation. Due to its nature, the cognitive gaps grow along with patients, and cognitive learning slows down, undergoing a 'progressive postnatal degenerative process'.[11]

Other structural abnormalities include an immature gyral pattern,[12] which inhibits neural processing;[13] irregular neocortical laminar formation,[14] essential for higher cognitive function; and delayed myelination of cortical fibers,[15] the neural pathways in the cerebral cortex responsible for a multitude of cognitive and sensory functions.[16] While these changes are hard to control or slow down after a certain age, intervention plans targeted at slowing down or stabilizing cognitive and socio-emotional degeneration at an early, developmental stage have a significant positive impact on people with Down Syndrome's lives.

Early Intervention plans aimed at improving cognitive and socio-emotional development:

Early intervention plans have higher success rates as compared to therapies conducted for older individuals with Down Syndrome,[17] as they take advantage of the high neuroplasticity of infant brains, intervening during periods of crucial, foundational, cognitive and socio-emotional (as well as overall) development.[18] Responsive Teaching, or RT, is one example of an early intervention plan. The program trains parents/caretakers of the patient on how to navigate interpersonal interactions with the patient to facilitate optimum developmental progress with more ease. The introduction of responsive interaction strategies into everyday correspondence encourages active engagement and learning processes, supporting more beneficial socio-emotional and cognitive development. There are 5 responsive interaction strategies. Reciprocity (using mirroring to join in an activity, being physically available and interactive, sustaining repetitive play, and maintaining face to face play void of external distractions such as toys), Contingency (taking the child's perspective, rapid identification and response to non-verbal requests, translating the patients actions into words, acknowledging unintentional vocalisations), Shared control (communicating without asking questions, imitating the patient's actions, giving the patient frequent opportunities to make independent choices, waiting for more mature responses rather than accepting every immature response from the patient), Affect (being animated, accompanying verbal communication with non-verbal communication, talking about novel or valuable actions of the patient), and Match (adjusting requests and interactions to the patient's unique developmental level). In a study conducted throughout the duration of a year, with 50 children with Down Syndrome and other similar developmental disorders, the developmental rate was proven to increase significantly with large jumps in the patients' capacities for expressive and receptive language.[19] Specifically, there was a 64% increase in cognitive development, 167% increase in expressive language, and 138% increase in receptive language, resulting in a 123% overall increase in developmental rate.

Another example of an early intervention plan is an early intervention preschool program. One empirical study investigating the early intervention focused on 4 key developmental areas in infants with Down Syndrome - Gross motor development, Fine motor development, Socio-emotional development, and cognitive and language development. The intervention plan implemented activities that encouraged movements such as turning, crawling, sitting, and walking, tailored to each infant's unique capabilities and progress, to support gross motor development. Activities centered around hand-eye coordination and manipulation skills were implemented to strengthen fine motor development. Instruction supported family activities and parent/caretaker involvement in aid of social and emotional development, and encouragement of communication through activities such as singing, finger games, and responsive talking was implemented to support cognitive and language development. Another study targeted and evaluated six key skill areas - gross and fine motor, cognition, receptive and expressive communication, and social self-help. Similar to the first study, parent involvement and instruction in all areas of instruction played a critical role in supporting the development of the children. Both studies display significant improvements in overall development. Hines and Bennett's study employed the Gesell Infant Development Scale (which assesses development in four areas - motor, adaptive, language, and personal-social skills) and assessed the infants with Down Syndrome at 4 - 6 months, at one year, and then at regular annual intervals. Tracking gross and fine motor skills, cognitive and language abilities, adaptive and interpersonal behavior, DQ(developmental quotient) of participants of the intervention plan were proven to remain stable, contrary to the trademark decrease in cognitive and socio-emotional abilities usually seen in patients of Down Syndrome, and directly in contradiction to the biological development of Down Syndromic brains.[4] Fewell and Oelwein's study assessed the DQ of its participants as well, but through calculation of the IR (Intervention Rate, the rate of develop-

DOI: 10.36838/IJHSR820.3

ment during the program) and the DRP(Developmental Rate Pretest - the rate of development observed throughout the initial years of the infant before implementation of the intervention plan) in proportion, factors which were calculated with respect to the average change of cognitive and socio-emotional development per month. Measured in percentages, the DRPs across the six developmental areas ranged from 0.58 to 0.68, while the IRs displayed significant increases ranging from 0.92 to 1.23. Cognitive development displayed an increase of 0.55, and developmental areas concerning socio-emotional development, such as receptive and expressive communication and social self-help (all of which constitute social awareness and interpersonal interactions), showed increases of 0.48, 0.56, and 0.43, respectively.[20]

The results of these studies show a largely positive, long-term effect on cognitive and socio-emotional development. Responsive Teaching tracked over a year demonstrates significant increases in the developmental rate of cognitive abilities and socio-emotional skills. The Early Intervention Preschool Program proves strong short-term improvements in important cognitive abilities, and socio-emotional abilities such as social awareness, interpersonal interactions, as well as a stable increase in the rate of long-term development in the initial years following the Early Intervention program. While these interventions considerably improve the quality of life for children with Down Syndrome by improving cognitive and socio-emotional development and therefore supporting a successful personal and social life, young adults with Down Syndrome characterize 'well-being' after rehabilitation largely as success in four domains of their lives – Relationships, Community participation, Independence, and Hopes for the future.[21] Improvement in skills such as receptive and expressive communication indicates an increase in the patient's interpersonal skills, and consequently their capacity to successfully navigate Relationships. This, as well as a marked improvement in social self-help, contributes to ease in community participation and independence, and an improvement in motor skills and cognitive development supports all four domains of a 'good life'. In this manner, the early intervention programs discussed, Responsive Teaching and the Early Intervention preschool program, significantly contribute to the well-being of Down Syndrome patients through childhood and young adulthood.

Conclusion

This literature review posits that early intervention plans have beneficial results towards improving cognitive and socio-emotional development prospects for people with Down Syndrome due to its congenital nature. Down syndrome is caused by Trisomy 21, a genetic mutation,[3] and can be characterized by trademark symptoms of dendritic atrophy, which correlates with cognitive and socio-emotional development rates slowing down as Down Syndrome progresses. Cognitive and socio-emotional rate of development is significantly lower in children with DS, and usually stagnates after 6-10 years of age.[10,18] This is supported by the structural differentiations observed, mainly the dendritic atrophy and the smaller shape of the brain.[10,11,13] Early intervention programs, such as preschool programs and responsive teaching, target maintaining or improving cognitive and socio-emotional developmental rates through targeting the symptoms, including interpersonal interactions, social engagement, and awareness, etc., and are effective in increasing the rate and the peak levels of cognitive and socio-emotional development.[4,8,17–20] Longitudinal research on early intervention programs is not available in large quantities, and therefore, this domain requires more extensive research to be done to more reliably gauge the general effectiveness of early intervention programs through childhood and adulthood, across children with Down Syndrome across the spectrum of severity. Structural abnormalities associated with Down Syndrome pose a considerable opportunity to learn more about the developmental barriers and causes of the atypical development but are not explored in depth past dendritic atrophy and a basic understanding of the visually atypical elements of the brain; an aspect that therefore demands to be researched more in-depth to investigate its role in socio-emotional and cognitive development more thoroughly, and gradually lead to a sure command of the causes of Down Syndrome and effective early intervention plans. Research on social cognition is another vast gap in research about DS, due to an inaccurate perception that children with Down Syndrome have high socio-cognitive abilities due to high levels of sociability, which is not and should not be considered a true indicator of the necessary socio-cognitive abilities required to maintain a healthy, social life, and longitudinal and bidirectional relationships.5 Therefore, future research should focus on socio-cognitive challenges that people with Down Syndrome face, particularly emotional recognition and Theory of Mind.

Acknowledgments

I would like to thank Monika Rybak as my mentor for her constant support through insightful feedback and suggestions throughout the process.

References

1. Fink, G. B.; Madaus, W. K.; Walker, G. F. A Quantitative Study of the Face in Down's Syndrome. Am. J. Orthod. 1975, 67 (5), 540–553. https://doi.org/10.1016/0002-9416(75)90299-7.
2. Abukhaled, Y.; Hatab, K.; Awadhalla, M.; Hamdan, H. Understanding the Genetic Mechanisms and Cognitive Impairments in Down Syndrome: Towards a Holistic Approach. J. Neurol. 2024, 271 (1), 87–104. https://doi.org/10.1007/s00415-023-11890-0.
3. Capone, G. T. Down Syndrome: Genetic Insights and Thoughts on Early Intervention. Infants Young Child. 2004, 17 (1), 45–58. https://doi.org/10.1097/00001163-200401000-00007.
4. Hines, S.; Bennett, F. Effectiveness of Early Intervention for Children with Down Syndrome. Ment. Retard. Dev. Disabil. Res. Rev. 1996, 2 (2), 96–101. https://doi.org/10.1002/(SICI)1098-2779(1996)2:2%253C96::AID-MRDD7%253E3.0.CO;2-V.
5. Cebula, K. R.; Moore, D. G.; Wishart, J. G. Social Cognition in Children with Down's Syndrome: Challenges to Research and Theory Building. J. Intellect. Disabil. Res. 2010, 54 (2), 113–134. https://doi.org/10.1111/j.1365-2788.2009.01215.x.
6. Drigas, A. S.; Karyotaki, M.; Skianis, C. An Integrated Approach to Neuro-Development, Neuroplasticity and Cognitive Improvement. Int. J. Recent Contrib. Eng. Sci. IT IJES 2018, 6 (3), 4. https://doi.org/10.3991/ijes.v6i3.9034.

7. Baurain, C.; Nader-Grosbois, N. Theory of Mind, Socio-Emotional Problem-Solving, Socio-Emotional Regulation in Children with Intellectual Disability and in Typically Developing Children. J. Autism Dev. Disord. 2013, 43 (5), 1080–1097. https://doi.org/10.1007/s10803-012-1651-4.
8. Baurain, C.; Nader-Grosbois, N. Socio-Emotional Regulation in Children with Intellectual Disability and Typically Developing Children in Interactive Contexts. Alter 2012, 6 (2), 75–93. https://doi.org/10.1016/j.alter.2012.02.001.
9. McLaughlin, K. A.; Garrad, M. C.; Somerville, L. H. What Develops during Emotional Development? A Component Process Approach to Identifying Sources of Psychopathology Risk in Adolescence. Dialogues Clin. Neurosci. 2015, 17 (4), 403–410. https://doi.org/10.31887/DCNS.2015.17.4/kmclaughlin.
10. Becker, L. E.; Armstrong, D. L.; Chan, F. Dendritic Atrophy in Children with Down's Syndrome. Ann. Neurol. 1986, 20 (4), 520–526. https://doi.org/10.1002/ana.410200413.
11. Benavides,Piccione, R. On Dendrites in Down Syndrome and DS Murine Models: A Spiny Way to Learn. Prog. Neurobiol. 2004, 74 (2), 111–126. https://doi.org/10.1016/j.pneurobio.2004.08.001.
12. Kemper, T. L. Down Syndrome. In Normal and Altered States of Function; Peters, A., Jones, E. G., Eds.; Jones, E. G., Peters, A., Series Eds.; Cerebral Cortex; Springer US: Boston, MA, 1991; Vol. 9, pp 511–526. https://doi.org/10.1007/978-1-4615-6622-9_12.
13. White, T.; Su, S.; Schmidt, M.; Kao, C.-Y.; Sapiro, G. The Development of Gyrification in Childhood and Adolescence. Brain Cogn. 2010, 72 (1), 36–45. https://doi.org/10.1016/j.bandc.2009.10.009.
14. Golden, J. A.; Hyman, B. T. Development of the Superior Temporal Neocortex Is Anomalous in Trisomy 21. J. Neuropathol. Exp. Neurol. 1994, 53 (5), 513–520. https://doi.org/10.1097/00005072-199409000-00011.
15. Kida, E., W., K. E. Abnormal Neurogenesis and Synaptogenesis in Down Syndrome Brain. Dev. Brain Dysfunct. 1994, 7, 289–301.
16. Price, D. J.; Kennedy, H.; Dehay, C.; Zhou, L.; Mercier, M.; Jossin, Y.; Goffinet, A. M.; Tissir, F.; Blakey, D.; Molnár, Z. The Development of Cortical Connections. Eur. J. Neurosci. 2006, 23 (4), 910–920. https://doi.org/10.1111/j.1460-9568.2006.04620.x.
17. Katz, S. Early Intervention Program for Infants with Down's Syndrome. Br. J. Nurs. 2000, 9 (8), 486–492. https://doi.org/10.12968/bjon.2000.9.8.6305.
18. Fidler, D. J.; Nadel, L. Education and Children with Down Syndrome: Neuroscience, Development, and Intervention. Ment. Retard. Dev. Disabil. Res. Rev. 2007, 13 (3), 262–271. https://doi.org/10.1002/mrdd.20166.
19. Mahoney, G.; Perales, F.; Wiggers, B.; Bob Herman, B. Responsive Teaching: Early Intervention for Children with Down Syndrome and Other Disabilities. Syndr. Res. Pract. 2006, 11 (1), 18–28. https://doi.org/10.3104/perspectives.311.
20. Fewell, R. R.; Oelwein, P. L. Effective Early Intervention: Results From the Model Preschool Program for Children with Down Syndrome and Other Developmental Delays. Top. Early Child. Spec. Educ. 1991, 11 (1), 56–68. https://doi.org/10.1177/027112149101100107.
21. Scott, M.; Foley, K.-R.; Bourke, J.; Leonard, H.; Girdler, S. "I Have a Good Life": The Meaning of Well-Being from the Perspective of Young Adults with Down Syndrome. *Disabil. Rehabil.* 2014, 36 (15), 1290–1298. https://doi.org/10.3109/09638288.2013.854843.

Author

Akanksha Mohanty is an 11th grader at Oberoi International School JVLR. She is passionate about Neuroscience and hopes to pursue a major in Neuroscience and a research career in that domain.

DOI: 10.36838/IJHSR820.3

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

The Vulnerability of Indigenous Rights In The Era of Global Extractivism: The Case of The Gateway Pacific Terminal

Aliya Abdesheva

Specialized Gymnasium 199, 89 Abish Kekilbayuly Street, Bostandyk District, Almaty, A15K3P9, Kazakhstan; alijushaa08@gmail.com
Mentor: Dr. Tanja Weis

ABSTRACT: Indigenous lands host 80% of the world's biodiversity. However, their environmental stewardship—driven by their high awareness of nature and confirmed by the UN Declaration on the Rights of Indigenous Peoples (UNDRIP)—is often viewed as an obstacle for industrial projects and global market access. While the conflict between sovereignty and globalization is well-documented, a significant gap remains in understanding how Investor-State Dispute Settlement (ISDS) mechanisms use data manipulation to bypass treaty-protected rights. This paper argues that, within ISDS arbitrations, private corporations use non-objective information, failing to meet standards of Free, Prior, and Informed Consent (FPIC) and violating Indigenous sovereignty. By analyzing the 1855 Treaty of Point Elliott, the USMCA, UNDRIP, the technical reports, and Joint Aquatic Resources Permit Applications (JARPA), this research examines the 2011–2016 conflict over the Gateway Pacific Terminal at Cherry Point (Xwe'chi'eXen). The findings suggest that corporations provide reconstructed data and non-objective environmental projections to subvert Indigenous opposition in favor of market access. This dysregulation is empowered by the ISDS mechanism, which allows foreign investors to sue Indigenous autonomies and host states in private tribunals if any protections interfere with expected profits. Ultimately, this paper demonstrates the vulnerability of Indigenous rights within modern trade regimes.
KEYWORDS: Human Rights, UNDRIP, ISDS, Lummi Nation, the Gateway Pacific Terminal.

Introduction

Only 6% of the world's population is Indigenous. All Indigenous peoples—ranging from the Inuit (North America) to the Māori (New Zealand)—occupy approximately 20% of the Earth's landmass, yet their territories preserve an estimated 80% of the world's remaining biodiversity.[1] This careful relationship with nature is inherent in Indigenous culture and lifestyle; formally, the ecological stewardship of Indigenous peoples is officially recognized and protected by international frameworks, such as the United Nations Declaration on the Rights of Indigenous Peoples (UNDRIP). UNDRIP protects individual and collective rights to self-determination, culture, identity, and land of Indigenous peoples. Moreover, introduced by UNDRIP, Free, Prior, and Informed Consent (FPIC) ensures that Indigenous communities have the right to give or withhold approval for projects on their lands, so that no work on their reservations can be started without the awareness and consent of Indigenous peoples. However, as global markets expand, Indigenous ancestral lands are not recognized as sovereign territories; instead, they are framed as obstacles to economic growth. International companies and private corporations, which aim to use Indigenous territories for their projects, operate through Investor-State Dispute Settlement (ISDS). This legal mechanism has threatened vulnerable Indigenous sovereignty, clashing environmental protection and Indigenous rights with corporate interests. Within the ISDS arbitration processes, investors present non-objective, reconstructed data that, in this case, was implemented to subvert Indigenous governance. In fact, ISDS lacks objective information in the context of Indigenous consent, allowing investors to disregard the standards of UNDRIP and FPIC. This demonstrates that FPIC and UNDRIP are "soft law" that fail to advocate for and protect Indigenous peoples.

A primary example of this struggle is found in the Salish Sea, an international waterway managed by over 65 Tribes and the home of First Nations—Coast Salish Sea People. Between 2011 and 2016, this region became the site of a critical jurisdictional dispute when project developer Pacific International Terminals (a subsidiary of SSA Marine) proposed the Gateway Pacific Terminal (GPT)—the largest coal export facility in North America—to be built on the Lummi Nation's sacred site, Cherry Point (Xwe'chi'eXen).

While the Lummi Nation's rights to withhold consent were promised by FPIC, SSA Marine pushed their project further with reconstructed environmental data and economic benefits, thus bypassing legal standards and Indigenous rights. This research argues that the eventual defeat of the GPT proposal was not achieved through the "soft law" of UNDRIP or FPIC, but rather through the domestic enforcement of the 1855 Treaty of Point Elliott. This "hard law" treaty provided a legal shield that resisted information manipulation within ISDS and preserved Lummi's land. Eventually, this dispute highlights a jurisdictional asymmetry, suggesting that Indigenous sovereignty remains in a globally vulnerable position unless substantive protections are integrated into the international trade regime itself.

In the first section, I present the work of modern scholars on the key concepts and legal frameworks central to this study. In the second section, I provide the synthesized methods used to demonstrate the violation of Indigenous rights within mod-

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.4

ern trade agreements. In the third section, I introduce the case study, establishing the historical, ecological, and cultural background of the Lummi Nation as well as the legal basis of their rights. Moreover, this third section analyzes the Gateway Pacific Terminal project, including its scale, environmental impact assessments, and the non-objective data provided by the developer. Finally, by examining the intervention of the high-level organization, I highlight the vulnerability and dependence of Indigenous rights within modern trade agreements.

Literature Review

Trade agreements and ISDS:

To cover the conflict between the Lummi Nation and the Gateway Pacific Terminal, it is first necessary to examine the evolution of the regional trade framework under which the dispute happened. The dispute over the terminal originally manifested under the 1994 North American Free Trade Agreement (NAFTA), which was subsequently replaced by the United States–Mexico–Canada Agreement (USMCA) on July 1, 2020. The USMCA was designed to transform the regional economy, aiming to be "a mutually beneficial win for North American workers, farmers, ranchers, and businesses."[2] The new agreement promised to create "more balanced, reciprocal trade supporting high-paying jobs for Americans and grow the North American economy."[2] While it boosted the North American trade economy, it also narrowed the "regulatory policy space" for governments to fulfill environmental and Indigenous treaty obligations.[3] To assess the broader legal framework underlying these disputes, it is necessary to examine the architecture of international arbitration. The structural mechanism driving these narrow regulation policies is the investor-State dispute settlement system (ISDS).

ISDS Criticism and Regulatory Chill:

The following section reviews existing scholarship on the structural biases of the ISDS framework and the phenomenon of "regulatory chill." Legal scholars, such as Barnali Choudhury, José E Alvarez, Sergio Puig, and Anton Strezhnev, have long been criticizing the ISDS system, arguing that it is controversial and biased.[4-6] During the ISDS arbitrations, investors can bring claims against government actions, even in domestic policy areas such as environmental protection and public health.[7] As Perrone notes, protections provided by the ISDS system are extended specifically to private corporations.[8] A primary consequence of the ISDS framework is the "regulatory chill"—a phenomenon in which governments refrain from enacting public interest policies due to the threat of costly arbitration or substantial compensation to foreign investors.[9] The consequences of ISDS within independent, specifically developing, states have already been meticulously studied. Specifically, case studies in the Global South—including environmental litigation in African countries, challenges against public health regulations in Latin America, and mining disputes in Central Asian countries—demonstrate that ISDS favors capital-exporting states.[10-13] However, little research has been conducted on how the ISDS mechanism impacts vulnerable Indigenous autonomies that operate under different legal protections than those of independent states. Moreover, there is a conceptual gap in studying methods used by investors, such as the utilization of reconstructed data and the violation of domestic law, to bypass Indigenous sovereignty and reframe their treaty-based status as right-holders to economic stakeholders. While scholars like Wang, Ning, and Zhang have begun studying how investors bypass local law, they have not yet applied it to tribal treaty rights in developed nations like the U.S.[14]

By ignoring Indigenous treaty rights, human rights, and environmental threats, the conflict within ISDS becomes "two-dimensional." This systematic limitation of disputes within ISDS is demonstrated by Lorenzo Cotula; in his work, he compares case studies of indigenous and agrarian groups, demonstrating that their human rights have been appropriated by social actors who reclaim their land, resources, and territories.[15] As Lorenzo Cotula demonstrates, ISDS excludes local communities, framing the disagreement as a binary struggle between a Host State and an Investor over a financial asset.[16] This challenge of exclusion is further analyzed by Gus Van Harten.[17] Van Harten's work supports the claim of one-sided wins, exploring how the public law system of ISDS allows corporations to challenge sovereign acts, like environmental protections and, in this case, treaty rights, in a private forum. As the environmental consequences and minority rights are often neglected within the ISDS arbitrations and placed in a private, commercial arbitration setting, host states tend to be "chilled" on regulatory policy and environmental protections.[17] The primary consequence of this dysregulation is the "regulatory chill". As Soloway observes, in international investment law, the phenomenon of the "regulatory chill" occurs when a government refrains from enacting or enforcing public interest policies, such as environmental protections, due to potential disputes with foreign investors.[18] Consequently, policymakers, such as the U.S. government, prioritize avoiding disputes and costly claims "over the development of efficient regulation in the public interest."[19] This asymmetry in regulation policies directly undermines Indigenous rights and challenges their environmental stewardship.

International protections for Indigenous rights:

The counterweight to the corporate pressure toward Indigenous communities is the United Nations Declaration on the Rights of Indigenous Peoples (UNDRIP). UNDRIP is more contemporary and global than the domestic 1855 Treaty, which will be mentioned and analyzed in the following sections. However, it does not introduce new or 'extra' rights to Indigenous peoples. In reality, its purpose is to guarantee Indigenous communities the same protections and access to justice as any other group.[20] To meet this goal, the UNDRIP has introduced Free, Prior, and Informed Consent (FPIC), providing the progressive right of self-determination for Indigenous Peoples.[20,21] It is important to acknowledge that the FPIC mechanism was introduced "exclusively for Indigenous and other 'traditional peoples', such as quilombolas in Brazil and tribal peoples in Africa", becoming a foundation for Indigenous peoples and local communities to shape the direction and consequences of resource interventions.[21,22] Specifically,

the goal of FPIC is to ensure that, giving or withholding their permission, Indigenous Peoples are consulted and cooperate in good faith with States before the implementation of any legislative or administrative measures that may affect their lands.[23] Article 32 highlights that "Indigenous peoples have the right to determine and develop priorities and strategies for the development or use of their lands or territories and other resources." In order to obtain Indigenous approval of any projects affecting their territories and resources, states shall consult with indigenous peoples through their own representative institutions. As stated in Article 32, Indigenous consent is highly required in projects connected to the development, utilization, or exploitation of mineral, water, or other resources.[24] Finally, Article 32 ensures states are responsible for the consequences of any projects developed on Indigenous land, affirming that "States shall provide effective mechanisms…for any such activities, and appropriate measures shall be taken to mitigate adverse environmental, economic, social, cultural or spiritual impact."

UNDRIP Limitations:

While UNDRIP has to protect Indigenous rights and advocate for their self-determination, it functions primarily as a non-binding moral framework. The limitations of the UNDRIP standards to protect Indigenous sovereignty are studied by Stavenhagen, who notes that Indigenous rights violations keep occurring, and the consequences of resource extraction on sovereign Indigenous lands remain a major problem for Indigenous peoples.[25] Further substantive work on the UNDRIP limitations has been done by Maro Barelli. Barelli argues that, while UNDRIP is a significant legal document, it remains "soft law."[26] Because it is a Declaration and not a formal Treaty, it does not create legally binding obligations for states in the same way a trade or investment agreement does. Moreover, with the rise of economic globalization, the responsibility to meet FPIC standards has moved from states to private corporations.[27] This created a significant implementation gap in the FPIC duties, reducing equal negotiations between Indigenous peoples and investors to a bureaucratic formality. The next phenomenon of the "governance gap" was mentioned, defining the failure of states' obligation to protect Indigenous rights.[21] This negligence is explained by a lack of "mechanisms to verify compliance with human rights responsibilities" and the absence of "legal instruments at the international level to enforce companies to comply."[21] This demonstrates a systematic failure in the implementation of human rights within Indigenous territories. Both private investors and states view UNDRIP standards as aspirational rather than mandatory, thus frequently circumventing critical responsibilities such as the FPIC process.

The further violation of Indigenous rights, UNDRIP, and FPIC within the ISDS system includes systematic data asymmetry. Corporate information manipulation has been thoroughly documented in traditional market contexts; specifically, Healy and Palepu analyze how firms maintain a technical monopoly by distorting financial statements and management disclosures.[28] However, little research has explored how the same information distortions within ISDS are applied to ecological data. In the context of the Gateway Pacific Terminal dispute, this asymmetry shifted from the financial to the environmental field. The project investors did not merely misrepresent the economic projections of their proposal; crucially, they presented reconstructed environmental risk assessments designed to minimize apparent impacts on local ecosystems, thus subverting the requirements of FPIC.

Methods

Through a single-case study of the conflict between the Lummi Nation and Pacific Terminals (PIT) over the proposed Gateway Pacific Terminal (GPT) at Cherry Point, this research aims to demonstrate the broad concept of Indigenous rights within modern trade agreements. The dispute over the GPT serves as a perfect critical case study as it intersects sovereign domestic Indigenous treaties, international trade agreements, globally recognized "soft law" frameworks, and corporate environmental reporting. Ultimately, this paper examines micro-mechanisms of corporate behavior and legal navigation within modern trade regimes toward vulnerable communities that may be missed in a broad multi-case study.

This study is built on three primary legal frameworks to determine the extent to which the territorial and human rights of Indigenous peoples are protected and respected within contemporary trade regimes. First, the study examines the power of the Treaty of Point Elliott (1855) with a specific focus on Article 5 regarding "Usual and Accustomed" fishing grounds, thus demonstrating how Lummi's territories and resources define tribal sovereignty and fishing rights within the Salish Sea. Besides, the paper analyzes the international "soft law" protections outlined in the UNDRIP and its core mechanism, FPIC. After legal sovereign protections, the research evaluates the specific environmental and logistical information submitted by the project developer within their official report of Joint Aquatic Resources Permit Application (JARPA). Subsequently, the study synthesizes official public interest letters from environmental advocacy groups like Salish Land Policy Solutions and the U.S. Army Corps of Engineers (USACE).

The analytical process was conducted in two stages. First, the legal basis for veto power under domestic treaty law was determined and compared with the consultation standards required by FPIC. Second, a cross-analysis of the data was conducted to identify information asymmetry. This was accomplished by comparing technical indicators provided by the developer (such as vessel traffic frequency and local pile driving zones) with biological and geographic impacts documented by independent federal and tribal agencies. By measuring the gap between corporate projections and empirical reality, the methodology allows us to identify "data reconstruction" strategies used to circumvent Indigenous sovereignty. To observe how these data strategies manifest in an actual regulatory dispute, this methodology must be attached to a specific geographic and ecological area. Consequently, the following section introduces the primary case site of this research, detailing how corporate industrial ambitions directly collided with the ancestral waters of the Lummi Nation.

DOI: 10.36838/IJHSR820.4

Case study

The Lummi Nation and The 1855 Treaty:

The Salish Sea is a complex international inland sea known for its rich and diverse marine life and ecosystem that contains 37 mammal species, 172 bird species, 261 fish species, and over 3,000 invertebrate species.[29] Being located in the Canadian province of British Columbia and the U.S. state of Washington, its waters are managed not only by the U.S. and Canadian governments but also by over 65 Tribes.[30] These tribes are considered to be the First Nations of the Salish Sea—Coast Salish peoples. In 1855, Coast Salish peoples signed with Washington Governor Issac Stevens the Treaty of Point Elliot.[31] In exchange for millions of acres of land, the U.S. government recognized and established independent reservations, including the Lummi Nation. The contemporary Lummi's estimated population is 5,000 tribal members, and its self-governing community manages around 13,000 acres of tidelands, including Bellingham Bay, the Strait of Juan de Fuca, parts of the Georgia Strait, and Cherry Point.[32] As "People of the Sea", the Lummi Nation relies on its waters: Lummi are famous for their reef-netting system, and their primary source of survival is salmon.[33] Cherry Point serves as a site of profound cultural and spiritual significance to the Lummi Nation, as this territory has been the ancestral homeland of their people for over 3,000 years.[34] As a result of the Point of Elliot Treaty—specifically Article 5 that guarantees all the Salish Sea tribes the right to fish at their "usual and accustomed grounds" (U&A), providing access to traditional food resources, tribal sovereignty, and resource management—Lummi Nation obtained perpetual rights to access and control their lands and resources. The 1855 Treaty of Point Elliot has been a legal shield for Coast Salish sovereignty for more than 160 years; however, in 2011, Indigenous rights protected by the treaty were challenged by private corporations and their multibillion-dollar proposal. In 2011, the project developer, Pacific International Terminals (a subsidiary of SSA Marine), proposed to construct the Gateway Pacific Terminal on the Lummi Nation's sacred site—Cherry Point (Xwe'chi'eXen).

The GPT proposal and Environmental Impact:

Before introducing the Gateway Pacific Terminal proposal, it must be acknowledged that, based on the employment report of Hodges, H., & Beyers, B., Cherry Point had already been used in heavy industry—specifically for large-scale fossil fuel refining, propane export, and aluminum smelting—for national and international markets.[35,36] Hosting major facilities like the BP refinery, Phillips 66 refinery, and former Alcoa plant, Cherry Point has experienced a severe decline in fisheries over the past half-century, with some populations facing total collapse.[37] Consequently, Cherry Point's vulnerable ecological position would have been threatened further by the construction of the massive coal terminal, and the proposed volume of extraction that will be mentioned in the following passage.

The Gateway Pacific Terminal (GPT) would have become the largest North American coal export facility. According to Pacific International Terminals, the project was designed to export up to 54 million metric tonnes of dry bulk to Asian markets annually, with coal accounting for 48 million tonnes of that total.[36] As the University of Washington Environmental Law and Regulation Practicum reports, the terminal construction would have excavated an estimated 3.4 million cubic yards of material—equivalent to over 680,000 dump truck loads—to grade the site for rail service, resulting in high levels of earth disturbance and increased soil vulnerability.[38,39] The Earth disturbance to construct the proposed project would have negatively impacted the rare biodiversity, local economy, and traditional life of the Salish Sea region. According to the JARPA filed by SSA, over 50 species of fish and larger organisms would have likely been affected. The proposed project encompasses 560 acres, 255 acres of which would have been directly disturbed. While some of the wetlands lost would have been attempted to be recreated on site, 140.6 acres were planned to be mitigated through means that had not been determined.[38] This severe disruption extends into the marine environment. Initial technical metrics framed the proposed 3,000-foot multi-berth wharf as standard marine infrastructure; however, independent empirical assessments revealed that driving hundreds of large-diameter steel piles into the coastal soil would permanently destroy herring spawning habitats, which are critical to the Salish Sea ecosystem (Table 1). Furthermore, data from BNSF Railway indicates that loaded railcars can lose up to 3% of their total mass as fugitive dust during transit.[40] Applied to the Gateway Pacific Terminal's maximum proposed volume of 48 million annual tons of Powder River Basin coal, a constant 3% transit loss would yield a projected atmospheric and terrestrial emission of approximately 1.44 million tons of fugitive coal dust per fiscal year.[38]

Data Asymmetry and Informational Gatekeeping:

In addition to the hazardous scale of the project, a further concern of this proposal was a lack of administrative transparency and rigorous data verification. According to Ehrlichman and Dykes, policy consultants have raised concerns regarding the "lack of good information and analysis available to the public" during the early stages of the GPT review process.[41] Specifically, Ehrlichman and Dykes' work indicated that SSA Marine's "factual premises about key planning numbers" were not supported by any documented evidence and lacked critical baseline information necessary to commence SEPA review. The lack of transparency effectively concealed the associated cross-border environmental risks inherent in the project's logistics. Specifically, the Washington State Department of Ecology's climate projections warned of intensifying precipitation and winter storm events; however, the project developer failed to account for the cumulative impact of stormwater leaching. The operation of 18 daily trains, each consisting of 125 coal cars, represents a significant, uncontrolled, non-point source of hazardous material discharge. Because active public rail lines border the Salish Sea coastline for extensive stretches and cross dozens of critical river and stream tributaries, this uncontrolled leaching of toxic substances threatens widespread water pollution—a vulnerability that national regulators have not addressed under standard Clean Water Act protocols.[41] Furthermore, the core operational parameters submitted by

the developer lacked any mandatory regulatory or voluntary limits. Although Pacific International Terminals claimed that the terminal's maximum capacity would be strictly limited to 18 trains per day and 48 million tons of exports per year, Erlichman and Dykes revealed that no such restrictive conditions or legal commitments were included in the JARPA applications or official project plans.[41] This structural omission allowed for potential, unregulated expansion driven by future improvements in the technological efficiency of the loading terminal. Finally, the applicant's projections relied on an unsubstantiated net-traffic assumption: that Canadian coal rail shipments would completely cease upon the GPT's completion, resulting in a net-zero increase in regional rail traffic. The developer did not provide any empirical data on current daily rail traffic volumes to Canada, thus failing to establish a verifiable historical baseline. The resulting lack of data has led independent policy advocates to demand that the Environmental Impact Assessment (EIA) report consider the proposed 18 daily trains not as a replacement, but as a massive, additional environmental burden on the regional ecosystem.[41]

Ergo, SSA Marine practiced informational gatekeeping, avoiding Indigenous concerns and ensuring that the regulatory process began on terms favorable to the investor. This regulatory deficit illustrates the specific patterns of corporate behavior examined by Van Harten.[17] This strategy forced the Lummi Nation into a reactive position, as tribes had to use their own limited resources to challenge and correct the initially reshaped data. Furthermore, this control over the initial data narrative represents a severe violation of FPIC. For consent to be "informed," the information provided must be transparent, objective, and accessible. However, the 2014 Vessel Traffic and Risk Assessment Study, conducted by The Glosten Associates for SSA Marine (Pacific International Terminals), contravenes the FPIC standards, containing further data reconstruction.[42] For instance, SSA Marine vice president Bob Watters argued that the study showed the terminal would have a "less than 1 percent impact" on treaty fishing. However, the same report contained data indicating that intense shipping would actually increase disruptions to Lummi fishing by 76%.[42-44] This finding shows the statistical erasure that was practiced by SSA Marine. By making the legal violation of Indigenous rights and environmental standards appear “de minimis,” or too small in a court of law, SSA Marine pushed their project proposal forward. This broader strategy of statistical erasure is demonstrated in Table 1, which contrasts the developer's self-reported metrics against the empirical findings verified by independent tribal and federal agencies.

Subsequently, the lack of data integrity was maintained through the exclusion of Indigenous peoples from important negotiations. As Lummi’s land was at the center of the dispute, the Lummi Nation, according to FPIC standards, had to be informed first about any project developments. However, based on the Salish Land Policy Solution report, the Office of Regulatory Assistance (ORA) MAP Team process excluded avenues for factual input from affected cities and counties along the transit line.[41] Specifically, Faith Lumsden, Director at ORA, had denied impacted jurisdictions, Bellingham, Skagit County, Burlington, and others, a seat at the ORA table. Excluding affected Indigenous communities made it impossible to track regional environmental damage and verify data presented by the project developer. Without this shared and transparent data, true informed consent could not happen.

Table 1 Legend: Comparing corporate projections with empirical environmental data. The data reveal a pattern of information asymmetry—specifically through cumulative omissions, baseline deficits, frequency truncation, and proxy distortions—utilized to minimize the projected ecological impacts on Indigenous lands and marine habitats.

Environmental Impact/ Feature	JARPA Technical Metric (SSA Marine Projections)	Empirical Impact Assessment (USACE, Tribal & Salish Sea Policy Findings)	Mechanism of Information Asymmetry
Vessel Traffic & Terminal Frequency	The projected Capesize bulk carrier volume is 487 vessels per year; the project is framed as an incremental, low-impact addition to existing regional maritime transport routes.	A minimum of 974 transits (one-way movements) would pass directly through treaty-protected *Usual & Accustomed* (U&A) fishing grounds, creating severe physical gear conflicts.	The omission of cumulative impacts. By modifying models to isolate terminal vessel traffic from regional congestion, SSA Marine hid the true cumulative risk of local vessel collisions.
Wharf & Structural Footprint Impact	A trestle and multiple berth structure, which extends approximately 3,000 feet into the water, is classified as standard deepwater infrastructure.	The construction requires driving hundreds of large-diameter steel piles into the coastal soil, which will permanently destroy vital herring spawning areas.	Baseline Data Deficit. By using global regional averages to calculate local mortality of benthic organisms (seafloor), the project developer masked the destruction of the unique herring population at Cherry Point.
Acoustic & Biophysical Disturbance	The industrial acoustic impact zone is localized directly on the industrial coastline and is maintained within standard regulatory threshold values.	Intense underwater acoustic propagation, covering long distances in narrow straits, will disrupt the echolocation of southern populations of killer whales.	Frequency Truncation: By utilizing atmospheric noise propagation models instead of marine hydroacoustic vectors, the project developer significantly underestimated the sound intensity underwater.
Overwater Shading & Primary Production	The shading area above the water surface is minimal, and the deck design minimizes continuous light blockage above the marine environment.	Continuous overwater shading across multiple acres of sensitive nearshore habitat will block sunlight and halt the primary production of eelgrass beds.	Proxy Metric Distortion. Replacing local light penetration patterns with macro-scale vegetation data obscures localized marine food web collapse.

Economic Benefits over Environmental Risks:

However, minimizing environmental impact and excluding locals was not enough to win in arbitrations. The greater strategy to obtain permission for construction, especially within the ISDS system, is to demonstrate economic benefits and projected profits. The construction of the GPT promised to stabilize the economy of Whatcom County and Washington State by

DOI: 10.36838/IJHSR820.4

creating 430 permanent jobs and by generating approximately $11 million in annual state and local tax revenues.[36] Being prioritized over environmental risks, economic projections served as a primary instrument for SSA Marine to push their project further through the U.S. domestic permitting process, directly challenging the U.S. federal government's legal obligation to uphold Lummi treaty protections. These numbers were provided to minimize and replace the potential hazards to human health and the environment. In this case, the emphasis on job creation serves as a distraction technique, attempting to justify the permanent ecological destruction of Cherry Point through temporary economic incentives. While jobs and tax revenues are subject to market fluctuations and have a limited trend, the destruction of the Pacific herring and Chinook salmon populations at Cherry Point is irreversible.

The conflict resolution and its factors:

The massive scale of this environmental damage and physical disturbance shows why the final legal decision was based on protecting treaty rights rather than taking economic promises into account. The conflict was resolved on May 9, 2016, by the United States Army Corps of Engineers (USACE) following the final and successful objection by the Lummi Nation. In cooperation with Pacific International Holdings, LLC, Lummi provided vast information regarding fishing practices, potential impacts, and mitigation to support their objective claims. After careful consideration of all the information available, including environmental assessment, the U.S. Army Corps of Engineers (2016) reported that it could not permit the terminal because the project's impacts on Lummi Nation treaty rights were found to be "greater than de minimis," and violated the Lummi Nation's U&A rights. According to the U.S. Army Corps of Engineers, "The district's evaluation of effects of the proposal on the Lummi's U&A fishing rights is undertaken to fulfill the federal government's responsibility to protect treaty rights.[45]

The Corps may not permit a project that abrogates treaty rights." While the five-year conflict over the GPT proposal ended in favor of the Lummi Nation, this fair and legal outcome was not guaranteed by the UNDRIP. As the case of the Belo Monte dam in Brazil—a large hydroelectric facility built without obtaining FPIC—shows, the "soft law" of the UNDRIP can be easily ignored by the government.[46] The mining activities of AngloGold Ashanti in the Canamomo Lomaprieta Indigenous Reserve, which were completed without the consent of the Indigenous communities, further demonstrate the lack of authority of the UNDRIP standards.[47] In the case of the Lummi Nation, the important factors of its success were the USACE intervention and the 1855 treaty. Firstly, the Lummi's inherent rights were effectively used by a federal agency that held the final power of interpretation. The Corps has chosen to honor its federal trust responsibility, thus prohibiting the project. However, had the USACE adopted a different definition of what constitutes a "de minimis" impact, or had political priorities shifted within the administration, the Lummi would have had a limited opportunity to stop the project. Secondly, the 1855 Treaty of Point Elliott is recognized under the U.S. Constitution as the "supreme law of the land".[48] It created a mandatory legal obligation for the federal government that could not be as easily bypassed as UNDRIP standards. The 1855 treaty has proven that Indigenous protection is most effective when it is grounded in domestic "hard law" that is non-negotiable to corporate expansion. Because of this chilled environment and systematic exclusion, the Lummi Nation had been fighting for its inherited rights and opposing the GPT for five years while its sovereignty was being silenced. By design, ISDS does not recognize the Lummi as a third sovereign state with an interest and rights in the land. Instead, the process treats the tribe's ancestral territory as part of state property or an investment site that the U.S. government is managing. This not only violates the legal protections, such as the 1855 treaty and the UNDRIP, but also ignores the sufferings and sacrifices, including the cession of 5 million acres of land, that Indigenous peoples have experienced to become independent.

■ Conclusion

Is the sovereignty of Indigenous peoples protected within modern trade agreements, and is this protection enough? While internationally recognized mechanisms are established, they fail to prevent human rights violations against Indigenous peoples.

As this research showed, the Lummi Nation's victory over corporate interests was achieved because of the exclusive "hard law" of the 1855 Treaty. The Treaty of Point Elliott could not be easily bypassed, as FPIC standards were. Additionally, the treaty found recognition and support among the U.S. Army Corps of Engineers. This big intervention is the second important factor of the Lummi's victory; it reveals the vulnerability of Indigenous rights, as not international organizations, but a high-level federal agency, was able to end this dispute. Nevertheless, the GPT conflict serves as a critical case study of how procedural formalism can be used to challenge tribal resources and defer sovereign justice to meet investor expectations. Future research should examine whether these findings hold in other global contexts, such as the Global South, where Indigenous communities may lack the "hard law" protections of historical treaties and must rely on "soft law" international frameworks. Furthermore, investigating how these legal vulnerabilities extend beyond land disputes to issues of cultural and intellectual property will be essential. Consequently, this case demonstrates that without robust domestic enforcement, the power of Indigenous sovereignty remains globally threatened; the GPT resolution may serve as a precedent for similar jurisdictional disputes across North American resource zones. Therefore, this paper highlights that until structural reforms are made to the international trade system, Indigenous sovereignty will remain dependent on the inconsistent protection of domestic courts.

■ Acknowledgments

I would like to express my gratitude to Dr. Ida Danewid for teaching me the very basics of research and for creating a motivating atmosphere. Special thanks to Dr. Tanja Weis for her

extensive mentorship and for providing the detailed feedback that was crucial in shaping this research.

This research was made possible by the IRIS Program at the Indigo Research Institute. I am deeply grateful for the full scholarship and technical training they provided, which taught me the skills necessary to conduct this study.

Finally, I would like to thank my family for their support and encouragement throughout the writing process.

References

1. United Nations Department of Economic and Social Affairs. Climate action. United Nations 2021. Available from: https://www.un.org/en/desa
2. Office of the United States Trade Representative. Countries & Regions. Executive Office of the President; 2026. Available from: https://ustr.gov/countries-regions
3. Labonté R, Crosbie E, Gleeson D, McNamara CL. USMCA (NAFTA 2.0): Tightening the constraints on the right to regulate for public health. Globalization and Health. 2019;15(1):35. Available from: https://doi.org/10.1186/s12992-019-0476-8
4. Choudhury B. Global Corporations and Sustainability: An Introduction. In: Choudhury B, editor. Global Corporations and Sustainability. Cheltenham, UK: Edward Elgar Publishing; 2025. Available from:https://digitalcommons.osgoode.yorku.ca/scholarly_works/3404
5. Alvarez JE. Are corporations "subjects" of international law? Santa Clara Journal of International Law. 2011;9(1):1–36. Available from: https://digitalcommons.law.scu.edu/scujil/vol9/iss1/1
6. Puig S, Strezhnev A. The David Effect and ISDS. European Journal of International Law. 2017;28(3):731–761. Available from: https://doi.org/10.2139/ssrn.2829006
7. Been V. NAFTA's investment protections and the division of authority for land use and environmental controls. Pace Environmental Law Review. 2003;20(1):19–43. Available from: https://doi.org/10.58948/0738-6206.1160
8. Perrone NM. Investment Treaties and the Legal Imagination: How Foreign Investors Play by Their Own Rules. Oxford, UK: Oxford University Press; 2021. Available from: https://doi.org/10.1093/oso/9780198862147.001.0001
9. Labonté R, Crosbie E, Gleeson D, McNamara CL. USMCA (NAFTA 2.0): Tightening the constraints on the right to regulate for public health. Globalization and Health. 2019;15(1):35. Available from: https://doi.org/10.1186/s12992-019-0476-8
10. Adekemi AO. Attracting African States' Participation in a Multilateral Investment Court. Cham, Switzerland: Springer Nature; 2024. Available from: https://doi.org/10.1007/978-3-031-73861-6
11. Fach Gómez K, Titi C. International Investment Law and ISDS: Mapping Contemporary Latin America. Journal of World Investment & Trade. 2016;17(4):515–535. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2784722
12. Caspian Policy Center. 2022 in Review: Key Events and Trends in the Greater Caspian Region. Washington, DC: Caspian Policy Center; 2022. Available from: https://api.caspianpolicy.org/media/ckeditor_media/2023/01/03/2022-review.pdf
13. Dumberry P. Requiem for Crimea: Why tribunals should have declined jurisdiction over the claims of Ukrainian investors against Russia under the Ukraine–Russia BIT. Journal of International Dispute Settlement. 2018;9(3):506–533. Available from: https://doi.org/10.1093/jnlids/idy022
14. Wang C, Ning J, Zhang X. International Investment and Indigenous Peoples' Environment: A Survey of ISDS Cases from 2000 to 2020. International Journal of Environmental Research and Public Health. 2021;18(15):7798. Available from: https://doi.org/10.3390/ijerph18157798
15. Cotula L. Between hope and critique: Human rights, social justice and re-imagining international law from the bottom up. Georgia Journal of International and Comparative Law. 2020;48(2):473–510. Available from: https://digitalcommons.law.uga.edu/gjicl/vol48/iss2/5
16. Cotula L. 'Land Grabbing' and International Investment Law: Toward a Global Reconfiguration of Property?. London, UK: International Institute for Environment and Development (IIED); 2016. Available from: https://www.iied.org/g04091
17. Van Harten G. The public-private distinction in the international arbitration of individual claims against the state. International & Comparative Law Quarterly. 2007;56(2):371–393. Available from: https://digitalcommons.osgoode.yorku.ca/scholarly_works/2340
18. Soloway J. NAFTA's Chapter 11: Investment protection, integration and the public interest. Choices. 2003;9(2):1–28. Available from: https://envireform.utoronto.ca/conference/nov2002/tollefson-soloway.pdf
19. Tienhaara K. Regulatory chill and the threat of arbitration: A view from political science. In: Brown C, Miles K, editors. Evolution in Investment Treaty Law and Arbitration. Cambridge, UK: Cambridge University Press; 2011. p. 606–627. Available from: https://researchportalplus.anu.edu.au/en/publications/regulatory-chill-and-the-threat-of-arbitration-a-view-from-politi-2/
20. Anaya J. The right of Indigenous peoples to self-determination in the post-declaration era. In: Charters C, Stavenhagen R, editors. Making the Declaration Work: The United Nations Declaration on the Rights of Indigenous Peoples. Copenhagen, Denmark: International Work Group for Indigenous Affairs (IWGIA); 2009. p. 184–199.
21. Hanna P, Vanclay F. Human rights, Indigenous peoples and the concept of Free, Prior and Informed Consent. Impact Assessment and Project Appraisal. 2013;31(2):146–157. Available from: https://doi.org/10.1080/14615517.2013.780373
22. Baez S. The "Right" REDD Framework: National Laws that Best Protect Indigenous Rights in a Global REDD Regime. Fordham Law Review. 2011;80(2):821–864. Available from: https://ir.lawnet.fordham.edu/flr/vol80/iss2/14/
23. Office of the United Nations High Commissioner for Human Rights (OHCHR). International trade and investment. Geneva, Switzerland: OHCHR; [updated 2026; cited 2026 Jan 10]. Available from: https://www.ohchr.org/en/topic/international-trade-and-investment
24. United Nations. United Nations Declaration on the Rights of Indigenous Peoples: resolution / adopted by the General Assembly. New York, NY: United Nations; 2007. Resolution No.: A/RES/61/295. Available from: https://www.un.org/development/desa/indigenouspeoples/wp-content/uploads/sites/19/2018/11/UNDRIP_E_web.pdf
25. Stavenhagen R. Human rights and indigenous issues: Report of the Special Rapporteur on the situation of human rights and fundamental freedoms of indigenous people. Geneva, Switzerland: UN Commission on Human Rights; 2003. Document No.: E/CN.4/2003/90/Add.3.
26. Barelli M. Free, prior and informed consent in the aftermath of the UN Declaration on the Rights of Indigenous Peoples: Developments and challenges ahead. The International Journal of Human Rights. 2012;16(1):1–24. Available from: https://doi.org/10.1080/13642987.2011.597746
27. Alvarez JE. Are corporations "subjects" of international law? Santa Clara Journal of International Law. 2011;9(1):1–36. Available from: https://digitalcommons.law.scu.edu/scujil/vol9/iss1/1

DOI: 10.36838/IJHSR820.4

28. Healy PM, Palepu KG. Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. Journal of Accounting and Economics. 2001;31(1–3):405–440. Available from: https://doi.org/10.1016/S0165-4101(01)00018-0
29. Gaydos JK, Pearson SF. Bird and mammal species of the Salish Sea: a compilation. Northwestern Naturalist. 2011;92(1):49–65.
30. Salish Sea Institute; Western Washington University. Recreation, conservation, and protection of the Salish Sea region. Bellingham, WA: Western Washington University; 2026. Emerging Issues in the Salish Sea Series, Report No. 12. Available from: https://salishsea.wwu.edu/emerging-issues-salish-sea-series
31. Sankofa Impact. The Treaty of Point Elliott. Seattle, WA: Sankofa Impact; [updated 2024; cited 2026 Jan 10]. Available from: https://sankofaimpact.org/the-treaty-of-point-elliott/
32. Salish Sea Institute. About the Salish Sea. Bellingham, WA: Western Washington University; [updated 2025; cited 2026 Jan 9]. Available from: https://salishsea.wwu.edu/about-salish-sea
33. Bellingham Regional Chamber of Commerce. Lummi Nation. Bellingham, WA: Bellingham Regional Chamber of Commerce; [updated 2025; cited 2026 Jan 9]. Available from: https://bellingham.com/lummi-nation/
34. Bellingham Unitarian Fellowship (BUF). Indigenous Rights and the Sacredness of the Salish Sea. Bellingham, WA: BUF; 2016 [cited 2026 Jan 10].
35. Hodges H, Beyers B. Employment at Cherry Point. Bellingham, WA: Center for Economic and Business Research, Western Washington University; 2014. Available from: https://cbe.wwu.edu/files/2021-11/Employment%20at%20Cherry%20Point%202014.pdf
36. Pacific International Terminal, Inc., Washington State: Joint Aquatic Resource Permit Application; 2010–2011.
37. Ury E. County Council moves to protect Cherry Point: What's next?. Bellingham, WA: RE Sources; 2019 Apr 9 [cited 2026 May 29]. Available from: https://www.re-sources.org/2019/04/cherry-point-background/
38. Nakamura D, Van Deren E, Marhofer E, Jablonski P, O'Neill S. Environmental impact analysis for the Gateway Pacific Terminal: Cherry Point, Washington. Seattle, WA: University of Washington School of Law, Environmental Law and Regulation Practicum; 2011. Available from: https://static1.squarespace.com/static/5510d902e4b0838d895e37a0/t/5542b091e4b004fd41a6a611/1430433937320/CherryPointUW.pdf
39. SSA Marine; Pacific International Terminals, Inc., Washington State: Joint Aquatic Resources Permit Application (JARPA) Form for Gateway Pacific Terminal. Olympia, WA: Washington State Governor's Office for Regulatory Innovation and Assistance; 2011. Report No.: COE JARPA 20110913. Available from: https://www.ezview.wa.gov/Portals/_1357/images/default/20110913%20COE%20JARPA%20for%20GeoTech%20Clearing%20JARPA.pdf
40. BNSF Railway Company. Custer Spur rail line expansion and infrastructure upgrade proposal: Gateway Pacific Terminal project. Fort Worth, TX: BNSF Railway; 2011. Available from: https://www.whatcomcounty.us/DocumentCenter/View/2859/MPP-and-SHR-Revision-Applications-Accepted---June-10-2011-PDF
41. Ehrlichman T, Dykes B. Re: Gateway Pacific Terminal pre-scoping considerations. Letter to Ted Sturdevant (Director, Washington State Department of Ecology). Bellingham, WA: Salish Land Policy Solutions; 2011 Oct 20. Available from: https://www.ezview.wa.gov/Portals/_1357/images/default/Ehrlichman%20%20Dykes%20-%20ECTS%207993%20-%20INSYSTEM(1).pdf
42. Glosten Associates. Gateway Pacific Terminal vessel traffic and risk assessment study. Seattle, WA: Pacific International Terminals; 2014 Nov 4 [cited 2026 May 29]. Report No.: Glosten File No. 14013.01. Available from: http://web.archive.org/web/20161129183341/http://eisgatewaypacificwa.gov/sites/default /files/content/files/20141104_Vessel_Traffic_and_Risk_Assessment_StudyGlosten_small_0.pdf
43. Gemme A. Assessing energy justice: The case of Xwe'chi'eXen, Cherry Point [master's thesis]. Bellingham, WA: Western Washington University; 2021. Available from: https://cedar.wwu.edu/wwuet/1028
44. Schwartz, R. How Lummi Nation's treaty stands against big coal. Seattle, WA: Sightline Institute; 2016 May 11 [cited 2026 May 29]. Available from: https://www.sightline.org/2016/05/11/how-lummi-nations-treaty-stand-beat-big-coal/
45. U.S. Army Corps of Engineers, Seattle District. Army Corps halts Gateway Pacific Terminal permitting process [press release]. Seattle, WA: USACE; 2016 May 9 [cited 2026 May 29]. Available from: https://www.nws.usace.army.gl/Media/News-Releases/Article/754951/army-corps-halts-gateway-pacific-terminal-permitting-process/
46. Jaichand V, Sampaio AA. Dam and be damned: the adverse impacts of Belo Monte on indigenous peoples in Brazil. Human Rights Quarterly. 2013;35(2):408–447. Available from: https://dx-.doi.org/10.1353/hrq.2013.0023
47. Resguardo Indígena Cañamomo Lomaprieta (RICL). Public statement from RICL, Colombia: AngloGold Ashanti usurps indigenous territory. Riosucio, Caldas: RICL; 2008 Apr 22.
48. United States. US Constitution, Article VI, Clause 2 [The Supremacy Clause]. In: National Archives and Records Administration. The Constitution of the United States: A Transcription. Washington, DC: NARA; [cited 2026 May 10]. Available from: https://www.archives.gov/founding-docs/constitution-transcript#toc-article-vi-

Author

Aliya Abdesheva is a Kazakhstani high school student specializing in art. Driven by a passion for Anthropology and Indigenous Advocacy, she intends to study Social Sciences/Anthropology in the United States. Her research explores how modern legal systems impact the cultural and ecological rights of Indigenous nations.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Discovery of Small Molecules Targeting VEGF Using Virtual Screening as Potential Therapies for Vascular Malformations

Andrew Wi

Archbishop Mitty High School, 5000 Mitty Way, San Jose, CA, 95129, USA; andrew.jw.wi@gmail.com

ABSTRACT: Vascular malformations comprise a diverse group of congenital vascular disorders that often decrease the quality of life through symptoms of pain, hemorrhage, and organ dysfunction. The protein VEGF modulates vascular remodeling and angiogenesis, as well as vascular permeability and tissue repair, by binding to receptors; excessive VEGF signaling in the presence of vascular malformations often intensifies structural abnormalities. Targeting VEGF can therefore help mitigate the effects of vascular malformations and slow progression, yet VEGF inhibitors remain underexplored. Initial druggability tests were conducted with geometric, energetic, and machine learning based site prediction tools to find and evaluate several potential ligand-binding pockets. To further match compound candidates with the potential sites, Pharmit's pharmacore-based virtual screening tool was used to discover 15 possible drugs with satisfactory overlays on the 7 selected interactions of the VEGF protein. Each candidate virtually simulated molecular docking using SwissDock, and all 15 had satisfactory or better binding affinity and orientation. Prioritizing medical relevance, Lipinski's Rule of Five was applied using SwissADME to eliminate 7 candidates. Finally, ProTox-II's silico toxicity test showed that 5 of the 8 remaining drugs had tolerable levels of toxicity according to their predicted LD50 and toxicity class data. The findings shed light on 5 new VEGF inhibitors that demonstrate strong efficacy, availability, and safety. These drugs represent potential and progress in inhibiting vascular malformation growth and development as a new form of non-surgical therapy for such.

KEYWORDS: Chemistry, Medicinal Chemistry, Drug Discovery, VEGF, Vascular Malformation.

Introduction

Vascular malformations are congenital abnormalities in blood vessels that stem from a variety of rare diseases and disorders.[1-3] These abnormalities often form during fetal development but can sometimes form as a result of an injury during adolescence or pregnancy.[4,5] Genetic mutations, such as hereditary hemorrhagic telangiectasia (HHT), can also contribute to the likelihood of a vascular malformation.[6-8] The severity of such malformations can range from asymptomatic lesions to severe lumps causing pain, bleeding, or organ dysfunction.[9-11] Despite progress in classification and treatment, essential information, such as precise causes or ideal treatment, for vascular malformations continues to remain unknown.[10,12,13]

Current standard strategies used to address vascular malformations are limited and cannot serve as long-term solutions.[14-16] Traditional approaches like surgical excisions, sclerotherapy, and embolization are effective in localized control but fail to exhibit preventative care.[7,15] Drug therapies such as the mTOR inhibitor sirolimus are promising and are considered the most optimal treatment, but their potential side effects and prolonged treatment present issues in use.[1,14,17] The absence of any FDA-approved pharmaceuticals for the treatment of vascular malformations also poses an issue for patients.[13,18] Thus, there is an urgent need for the development and discovery of safer, long-term treatments with greater efficacy.[10,19]

To address such concerns, recent efforts to find drugs have centered around the identification and inhibition of relevant molecular pathways.[9,20] Inhibitors for the PI3K/AKT/mTOR and RAS/MAPK/ERK pathways are seen as those with the most potential, due to their importance in endothelial cell proliferation, migration, and differentiation.[3,21] PIK3CA inhibitor alpelisib and other MEK inhibitors have promising preclinical trial results that could reduce lesion size or alleviate symptoms.[19,22] However, the challenges of shifting scientific findings to accessible treatments with long-term safety remain and prevent such drugs from becoming mainstream therapies.

The signaling protein vascular endothelial growth factor (VEGF) regulates angiogenesis and is related to processes such as embryonic development. VEGF can trigger the proliferation and migration of new blood vessels through its binding to receptors such as VEGFR-1 and VEGFR-2.[15,22] While VEGFR-1 plays the role of regulating VEGFR-2, VEGFR-2 is the main receptor that drives endothelial cell angiogenesis. The main process related to VEGF remains angiogenesis, but its signaling also relates to vascular permeability and tissue repair.[4,17] With its diverse functions, VEGF remains relevant for physiological and pathological vascular processes.

Recent findings have shown a link between excessive VEGF signaling and the growth of vascular malformations.[21,22] Abnormally high VEGF signaling often disrupts normal vessel development and stability, which creates structurally anomalous vessels. Especially in AVMs, elevated VEGF levels can also exacerbate vascular malformations and foster further deviation.[5,12] The obstruction of VEGF signaling, on the other hand, using agents like bevacizumab has proven to decelerate lesion growth or even reverse early-stage vascular malformations.[3,21] VEGF signaling also often triggers other pathways related to vascular malformations, such as the PI3K/AKT/

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.5

mTOR and RAS/MAPK/ERK pathways, enhancing the effectiveness of VEGF drug targets as candidates for vascular malformation therapy.[9,19] The discoveries related to VEGF signaling and vascular malformation open up a new field of research for therapeutic strategies in addressing vascular malformations.[10,16,18]

The 2025 review "Genetic Insights Into Hemorrhagic Stroke and Vascular Malformations: Pathogenesis and Emerging Therapeutic Strategies" reveals VEGF's therapeutic potential in vascular malformation treatment.[23] The discovery of an intersection between RAS/MAPK/ERK pathways and VEGF signaling to trigger and manage function. There was also a strong link between VEGF drug candidacy and effective treatment for specifically AVMs and CCMs. The use of VEGF-inhibiting drugs such as bevacizumab was shown to normalize vascular permeability and to lessen lesion growth. Such efficiency was also extended to effectively address HHT through VEGF regulation.

Another 2025 study, "Extracranial arteriovenous malformations: towards etiology-based therapeutic management," focused on the treatment of AVMs by reducing VEGF overexpression.[24] Increased VEGF levels were shown to cause a greater risk of hemorrhage and vascular malformation relapse even after embolization. Preclinical settings demonstrated how bevacizumab reduced dysplasia, and early human clinical trials also yielded efficacy. Such promising findings prompted the BevacizuMAV phase II-III trials to begin to evaluate optimal dosing.

The recent 2025 paper "Vascular endothelial growth factor signaling in health and disease: from molecular mechanisms to therapeutic perspectives" focused on the phosphorylation of VEGFR2 to study how VEGF signals endothelial cell growth.[25] Specific phosphorylation sites on VEGFR2, such as Y951 or Y1214, were shown to be critical in processes like angiogenesis seen in vascular malformations. Mice with targeted mutations at specific phosphorylation sites were able to reduce vascular leakage and improve vessel stability, highlighting the possibility of precise interventions for VEGFR1/VEGFR2 signaling to treat vascular malformations.

Methods

Geometric-Based Binding Site Identification:

The ProteinPlus tool was accessed through the internet: ProteinPlus is a tool that identifies possible binding sites on different proteins based on their size and shape.[26] It mainly uses volume and surface area to determine how ideal the site is. To identify the VEGF protein, its PDB, a protein data bank code that is uniquely assigned to each discovered protein, of "4AF8" was typed into the blank text box. After clicking "go!", the "DoGSiteScorer Binding Site Detection" and then the "DoGSiteScorer" were clicked to access the actual geometric-based binding site identification tool. As there are no special complications, standard settings were set and then "Calculate" was clicked. To see each respective site's size, shape, location, and configuration in a different color, "show all pockets" was checked off.

Energetic-Based Binding Site Identification:

The FTSite tool was accessed through the internet: FTSite is a tool that identifies possible binding sites on different proteins based on their interactions with van der Waals probes.[27] To distinguish the identification of VEGF from other previous identifications, the job name was titled "VEGF." To identify the VEGF protein, its PDB ID of "4AF8" was typed into the blank text box. After clicking "Find My Binding Site", each binding site candidate was shown in a different color, but no additional information was given.

Machine Learning Site Identification:

The PrankWeb tool was accessed through the internet: PrankWeb is a tool that displays a molecular model of the protein with its possible binding sites.[28] It uses a variety of factors to score each site, but the amino acids making up each site can also be accessed. "Experimental Structure" was clicked to access the machine learning site identification tool. To identify the VEGF protein, its PDB of "4AF8" was typed into the "PDG" text box. After clicking "Submit", a model of the VEGF protein appeared, which could be moved and zoomed in to show specific details in the protein and its sites, as well as the amino acids that make up the protein and its sides.

Virtual Screening:

The Pharmit tool was accessed through the internet: Pharmit is a virtual screening tool that uses interactions in different proteins to match them with existing small-molecule drugs.[29] To identify the VEGF protein, its PDB of "4AF8" was typed into the "PDB" text box. After clicking "submit" and turning "Receptor Surface Opacity" to the lowest, a model of the interactions within the VEGF protein appeared. "Receptor Surface Opacity" is the opacity of the VEGF protein, and setting it to low ensured that the interactions are more visible. Clicking "ON" for all "Hydrogen Donor", "Hydrogen Acceptor", and "Hydrophobic" interactions balanced the the seven interactions are balanced between too many and too few so that there were a large amount of results, but also so that all the results had low enough RMSDs.

Molecular Docking:

The ChemDiv website was accessed through the internet to identify each small molecule. The each name small molecule ""X###-####" was typed into the blank text box. After clicking "enter" and clicking "compound profile", finding and copying "smiles" noted the SMILES, or the code that describes the structure of a compound or molecule in a way that is readable by online chemistry tools, for further processes. The SwissDock tool was accessed through the internet: SwissDock is a tool that virtually runs molecular docking to see if selected compounds are able to dock and bind to a protein target.[30] To identify the small molecule, its SMILES from was typed in the "Provide a SMILES" box, and "Prepare ligand" was clicked to prepare the small molecule. To identify the VEGF protein, its PDB of "4AF8" was typed in the "Provide a PDB ID" box, and "Prepare target" was clicked to prepare the protein. The "Define search space" was used to remodel the cross-section,

and trial and error was led to the initial values 22, 23, 32 for the "Search box center" box and the values 44, 46, 51 for the "Search box size" box. Clicking "Check parameters" provided an estimated time, and, depending on the small molecule, if the estimated time was over 1:00:00, going back to "Define search space" and shrinking the cross-section size until the estimated time was under 1:00:00 made the process more efficient. To be notified of the results after the listed time, email and docking name were entered before clicking "START DOCKING."

Drug Properties:

The SMILES for each small molecule was found and copied before accessing the SwissADME tool through the internet: SwissADME is a tool that finds drug properties to see if they follow Lipinski's rules and can serve as a satisfactory drug.[31] The previously copied SMILES was pasted in the "Enter a list of SMILES here:" box, and "Run!" was clicked.

Toxicity:

The SMILES for each small molecule was found and copied before accessing the ProTox-II tool through the internet: ProTox-II tool is a tool that identifies the level of toxicity in different drugs.[32] Clicking "TOX PREDICTION" led to a "Canonical Smiles" box where the previously copied SMILES was pasted, and clicking "smiles", "All", and "Start-Tox Prediction" began the process.

Results and Discussion

Geometric-Based Binding Site Identification:

Generally, a 300-700 Å^2 surface area is considered "druggable" for binding sites based on geometric analysis.[33] Smaller or larger sizes are often difficult to bind small molecules to efficiently. With surface areas smaller than 300 Å^2, a lack of contact points can obstruct proper binding.[34] On the other hand, surface areas greater than 700 Å^2 can lead to weak or nonspecific binding, as well as insufficient concavity to contain the drug.[35] Thus, when analyzing the 10 possible sites found based on geometric-based identification, G_P6, G_P7, G_P8, and G_P9 are all undruggable due to their less than 300Å^2 surface area (Table 1). While G_P0 is also larger than typical druggable sites, it remains the ideal candidate due to its location (Figure 1). As the site is a pocket with most of its sizable surface area being enclosed within the protein, concerns about concavity and nonspecific binding are better addressed.

Table 1: This table contains data representing the 10 binding sites found through geometric analysis. Each row lists the volume (Å^3), surface area (Å^2), color code, and druggability based on the previous factors. G_P0 is the most ideal binding site candidate.

	G_P0	G_P1	G_P2	G_P3	G_P4	G_P5	G_P6	G_P7	G_P8	G_P9
Volume (Å^3)	873.09	543.23	440.38	279.19	174.4	173.18	131.84	126.89	121.41	101.18
Surface Area (Å^2)	910.38	647.67	535.62	480.86	380.47	356.91	271.37	277.96	224.56	273.48
Drug Score	0.83	0.86	0.8	0.57	0.33	0.31	0.19	0.27	0.34	0.13
Color	Yellow	Purple	Neon Green	Red	Blue	Muddy Green	Pink	Teal	Orange	Indigo

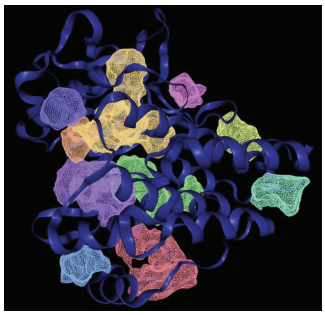


Figure 1: The protein plus site identified 10 different binding sites on the VEGF protein based on the size and location of each pocket. Each pocket is color-coded and has "druggability" evaluated based on accessibility and suitability for small molecules. From left to right and top to bottom, the colors of the pockets are: indigo, yellow, pink, orange, muddy green, purple, neon green, teal, blue, red. The yellow pocket shown in the figure is the most ideal binding site candidate.

Energetic-Based Binding Site Identification:

For a binding site to be "druggable" based on energy, sites are identified based on the levels of energy and interactions within pockets. Size and location are less important than the presence of sufficient interactions that allow for small molecules to reach the pockets and bind to the protein to initiate inhibition. In analyzing the resulting figure, three colored pockets near the center of the protein prove the protein is druggable (Figure 2).



Figure 2: The FTSite identified binding sites marked by colors, based on the energies of sites on the protein. All three colored pockets can serve as sufficient binding sites.

Machine Learning Site Identification:

In terms of the machine learning tools used by Prankweb to identify binding sites, a virtual model of the protein is analyzed to factor in location and amino acids making up pockets and assign "scores." Scores above 1.0 are considered satisfactory, and scores above 5.0 are considered ideal. As ML_P1, ML_P2, and ML_P3 are each above 1.0 (Table 2), VEGF displays satisfactory sites, and ML_P0's score of 38.53 provides the most ideal binding site for small molecules (Figure 3).

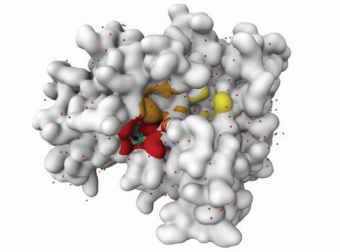


Figure 3: The Prankweb site identified four, color-coded binding sites based on machine learning, with key factors for each pocket being spatial distribution and amino acid composition. Each pocket was scored and ranked based on such factors. The red pocket is the most ideal binding site candidate.

DOI: 10.36838/IJHSR820.5

Table 2: This table lists each pocket's color, number of amino acids, and a few of the amino acids making up the pocket. The greatest comprehensive score ranks the four pockets based on the listed properties.

	ML_P0	ML_P1	ML_P2	ML_P3
Ranking	1	2	3	4
Score	38.53	1.55	1.44	1.06
Color	Red	Yellow	Orange	Blue
Number of Residues	37	7	8	6
Amino Acids	GLY117, LEU180, ASP191	SER249, ASP203, GLU242	ASP201, PHE192, ARG196	ARG172, SER231, MET217

Virtual Screening:

Using the tool Pharmit on a pharmacore map with 7 interaction spheres (Figure 4), an RMSD score is given to each compound that is virtually screened based on how well it layers with the protein. RMSD scores, or root mean square deviation scores, are, in simplest terms, calculations of the average distance between two sets of values. In the context of Pharmit and pharmacore mapping, RMSD scores are used to calculate how well compounds can fit a protein.[36] Such is done by aligning the two in the most ideal overlay and comparing distances between corresponding atoms (Figures 5-19). Beyond simple molecular structures, however, the pharmacore map uses existing interactions within the VEGF molecule to ensure that the small molecules are able to interact and potentially bind with the protein.

Scores between 0.5 and 0.7 would be considered acceptable. Of the 15 drug candidates that were virtually screened, only ChemDiv-4393-0009 has an RMSD score below 0.7 (Figure 5). However, since the other 14 candidates each have RMSD scores that have a deviation of less than 0.1, greater than the acceptable range of 0.5-0.7 (Table 3), the virtual screening can be considered successful for these candidates as well.

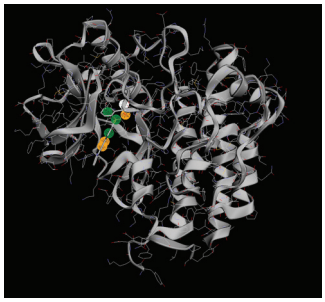


Figure 4: The protein VEGF is shown on a pharmacore map with 7 main interactions in different colors: white represents a hydrogen donor, yellow represents hydrogen acceptors, and green represents hydrophobic interactions.

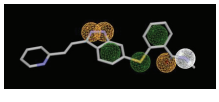


Figure 5: Compound 4393-000. Interacting with the target.

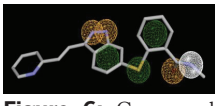


Figure 6: Compound V029-1633. Interacting with the target.

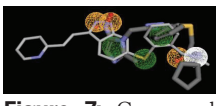


Figure 7: Compound 8017-4590. Interacting with the target.

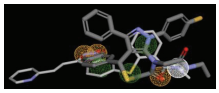


Figure 8: Compound V015-7427. Interacting with the target.

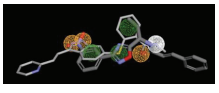


Figure 9: Compound V027-6087. Interacting with the target.

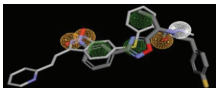


Figure 10: Compound V027-2157. Interacting with the target.

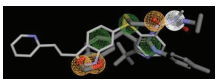


Figure 11: Compound V015-2773. Interacting with the target.

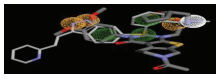


Figure 12: Compound K305-0525. Interacting with the target.

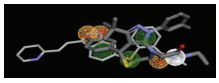


Figure 13: Compound V013-5298. Interacting with the target.

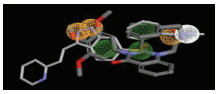


Figure 14: Compound C260-0238. Interacting with the target.

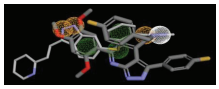


Figure 15: Compound G923-0287. Interacting with the target.

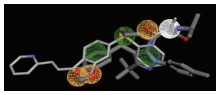


Figure 16: Compound V010-4661. Interacting with the target.

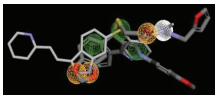


Figure 17: Compound V004-4933. Interacting with the target.

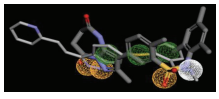


Figure 18: Compound L299-0508. Interacting with the target.

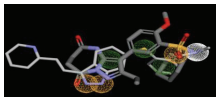


Figure 19: Compound L299-3577. Interacting with the target.

Table 3: The table lists each compound, its SMILES, corresponding figure, and data from the virtual screening. With small molecule mass and interaction, an RMSD score is given to evaluate how well each candidate interacts with the protein. ChemDiv-4393-0009 is the most ideal candidate, but all 14 are acceptable.

Compound	RMSD	Mass	RBnds	Figure	SMILES
ChemDiv-4393-0009	0.673	429	9	5	C1(C(=O)C2=CC=C(C=C2)F)C(C2=CC=CC=C2F)N(C(=O)C=1O)CCCCCC(O)=O
ChemDiv-V029-1633	0.704	415	8	6	CCC(C)C(c1nc(c2cccc(c2)[N+]([O-])=O)no1)NC(c1ccccc1[Cl])=O
ChemDiv-8017-4590	0.723	442	9	7	Cc1ncc([N+]([O-])=O)n1Cc1nnc(n1c1ccccc1)SCC(NC1CCCC1)=O
ChemDiv-V015-7427	0.745	557	7	8	C1CC1NC(CN1C(CSC(c2ccc3c(c2)OCO3)c2c(c3ccccc3)nn(c3ccc(cc3)F)c12)=O)=O
ChemDiv-V027-6087	0.749	392	8	9	CC(C(c1onc(c2cc([N+]([O-])=O)ccc2)n1)NC(/C=C/c3ccccc3)=O)C
ChemDiv-V027-2157	0.754	398	8	10	CC(C)C(c1nc(c2cccc(c2)[N+]([O-])=O)no1)NC(Cc1ccc(cc1)F)=O
ChemDiv-V015-2773	0.759	533	7	11	Cc1ccc(cc1)n1c2c(C(c3ccc4c(c3)OCO4)SCC(N2CC(NC2CC2)=O)=O)c(C(C)(C)C)n1
ChemDiv-K305-0525	0.759	543	9	12	CC(N1CCc2c3C(N(C(N(CC(NCC4CCCO4)=O)c3sc2C1)=O)c1ccc(c(c1)OC)OC)=O)=O
ChemDiv-V013-5298	0.762	533	7	13	Cc1cccc(c1)n1c2c(C(c3ccc4c(c3)OCO4)SCC(N2CC(NC2CC2)=O)=O)c(C(C)(C)C)n1
ChemDiv-C260-0238	0.762	475	9	14	COc1cc(cc(c1OC)OC)N1C(c2ccccc2N(CC(NCc2ccccc2)=O)C1=O)=O
ChemDiv-G923-0287	0.763	531	9	15	COc1cc(c(cc1c1cc(C(NCc2ccc(cc2)F)=O)c2c(n1)n[nH]c2c1ccc(cc1)F)OC)OC
ChemDiv-V010-4661	0.764	535	7	16	CC(C)NC(CN1C(CSC(c2ccc3c(c2)OCO3)c2c(C(C)(C)C)nn(c3ccc(C)cc3)c12)=O)=O
ChemDiv-V004-4933	0.771	609	9	17	COc1ccc(cc1)n1c2c(C(c3ccc4c(c3)OCO4)SCC(N2CC(NCc2cccc2)=O)=O)c(c2ccccc2)n1
ChemDiv-L299-0508	0.775	467	4	18	Cc1cc(C)c(c(C)c1)NS(c1cc(ccc1C)c1c(C)nn2C(CCC(Nc12)=O)=O)(=O)=O
ChemDiv-L299-3577	0.776	461	6	19	CCc1c(c2ccc(c(c2)S(NC2CCCCC2)(=O)=O)OC)c2NC(CCC(n2n1)=O)=O

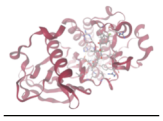
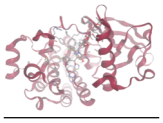
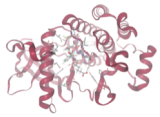
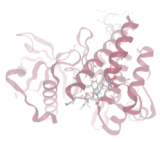
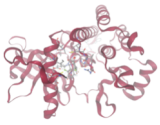
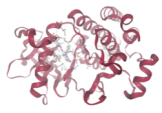
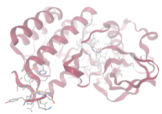
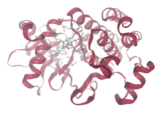
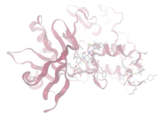
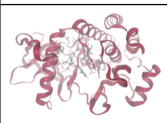
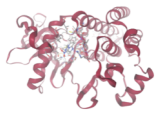
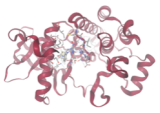
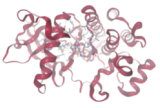
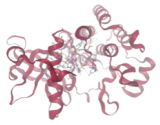
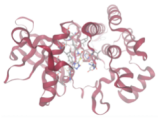
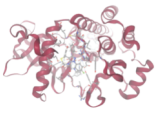
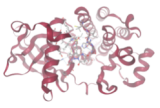
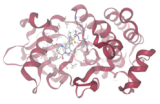
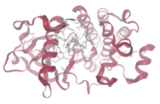
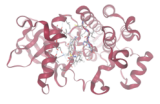
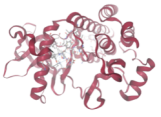
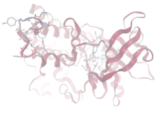
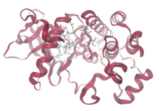
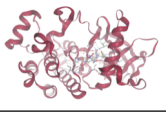
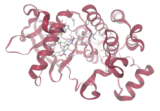
Molecular Docking:

Molecular docking is one of the most important experiments in a drug discovery project. Using a virtual model of the protein and compound, realistic simulations are run between the drug and every possible docking point and approach on the VEGF protein's surface. Testing docking evaluates how well a compound can approach and bind to the protein in order to begin any sort of inhibition.[37] The results from this experiment portray the true feasibility of small molecule drugs in acting as a possible therapeutic for vascular malformations.

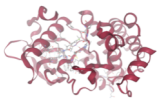
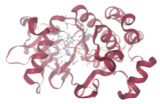
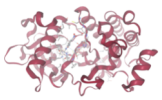
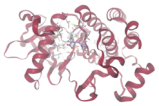
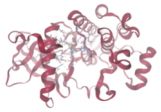
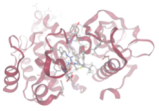
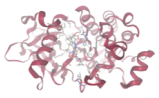
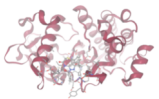
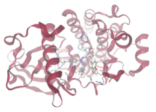
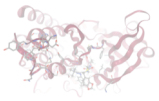
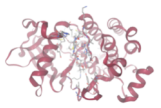
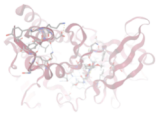
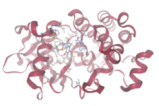
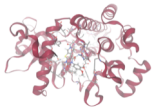
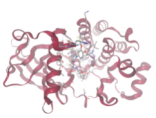
To determine the spontaneity of docking each drug candidate, the SwissDock tool was used to simulate molecular docking across the entire VEGF protein's surface. The parameters for each molecular docking simulation were set after

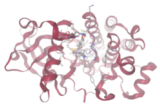
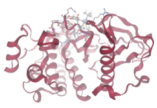
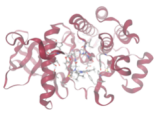
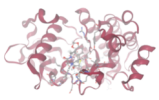
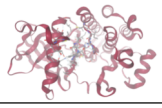
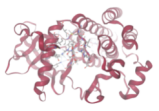
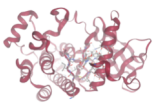
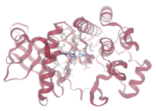
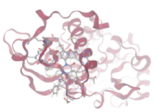
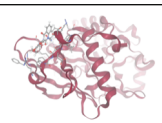
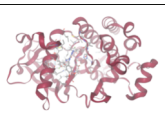
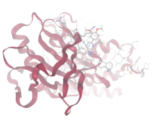
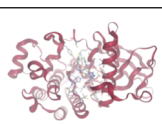
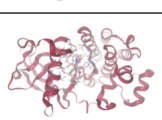
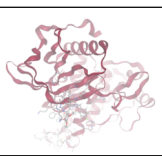
various trials and adjustments to determine the minimum parameters that would account for the maximum surface area. Each candidate's attempt at docking was given a SwissParam score, with an absolute value of 7-8 being satisfactory, 9 being great, and 10 being ideal (Tables 4-18).[38] As the compound ChemDiv-V027-6087 had a SwissParam score of -10.0058 (Table 8), it is considered an ideal candidate. Additionally, there were also 4 candidates that had SwissParam scores of -9, which provided great results as well (Table 6, Table 9, Table 13, Table 14). The 10 other candidates each had scores ranging from -7 to -8, with none having an absolute value less than 7, so all 15 compounds can be considered for the next phase of research.

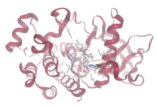
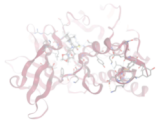
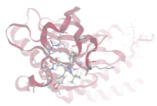
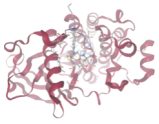
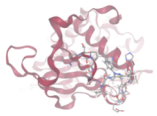
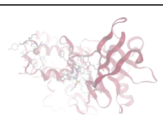
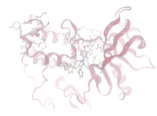
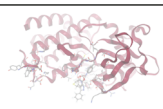
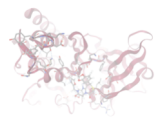
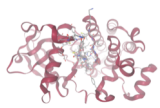
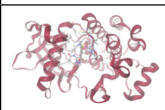
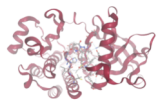
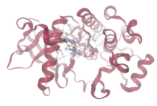
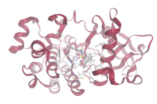
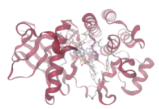
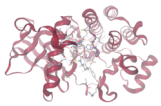
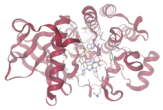
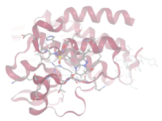
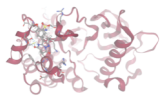
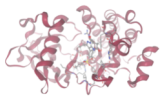
Table 4: Molecular docking of compounds on the VEGF protein.

Compound	Result	Cluster Number	SwissParam Score	Binding Pose of Compound to VEGF
ChemDiv-4393-0009	Satisfactory	0	-7.4869	
		3	-7.3202	
		5	-7.4940	
		6	-7.8000	
		7	-7.4028	
ChemDiv-V029-1633	Satisfactory	5	-8.8704	
		3	-8.4975	
		4	-8.3086	
		6	-7.8442	
		1	-7.7287	
ChemDiv-8017-4590	Great	1	-9.1069	
		4	-8.8363	
		0	-8.7276	
		3	-8.6509	
		6	-8.0580	
ChemDiv-V015-7427	Satisfactory	0	-8.2343	
		1	-8.2016	
		2	-8.1032	
		4	-7.9576	
		0	-7.9241	
ChemDiv-V027-6087	Ideal	1	-10.0058	
		2	-9.4998	
		0	-9.0186	
		3	-8.8561	
		4	-8.7069	

DOI: 10.36838/IJHSR820.5

ChemDiv-V027-2157	Great	0	-9.7904	
		2	-9.1912	
		1	-8.7009	
		6	-8.6618	
		3	-8.5877	
ChemDiv-V015-2773	Satisfactory	0	-8.2597	
		2	-8.2429	
		2	-7.9334	
		3	-7.5495	
		5	-7.4495	
ChemDiv-K305-0525	Satisfactory	1	-8.5544	
		5	-8.3300	
		3	-8.2011	
		8	-7.8659	
		0	-7.8548	

ChemDiv-V013-5298	Satisfactory	0	-8.1580	
		4	-7.4057	
		0	-7.3134	
		8	-7.3049	
		1	-7.2493	
ChemDiv-C260-0238	Great	0	-9.1128	
		3	-8.9344	
		4	-7.8125	
		1	-7.7466	
		5	-7.4370	
ChemDiv-G923-0287	Great	2	-9.4199	
		1	-8.7616	
		0	-8.7069	
		4	-8.4481	
		3	-8.3458	

ChemDiv-V010-4661	Satisfactory	8	-7.9318	
		0	-7.4391	
		7	-7.3112	
		6	-7.1171	
		5	-7.0050	
ChemDiv-V004-4933	Satisfactory	0	-8.3058	
		1	-8.0859	
		4	-7.9022	
		2	-7.8650	
		8	-7.8318	
ChemDiv-L299-0508	Satisfactory	1	-8.2658	
		0	-8.2153	
		6	-7.7609	
		2	-7.7220	
		7	-7.6990	
ChemDiv-L299-3577	Satisfactory	0	-8.0102	
		2	-7.6432	
		6	-7.2487	
		8	-7.2319	
		9	-7.1352	

Drug Properties:

While structure and compatibility are important for binding with proteins, small molecules must first permeate through membranes and cycle out of the body system after their work is done. Factors such as hydrogen bonds, mass, and solubility each affect how permeable compounds are, according to observations made by Christopher A. Lipinski in his Rule of 5.[39] As his conditions are maximums, any compounds with qualities less than his standards are more likely to be permeable and absorbable. Thus, the 8 qualified candidates are proven fully adequate in all aspects, and the only trial that remains is to evaluate their toxicity.

SwissADME tests compounds for absorption to ensure that drugs can reach their protein targets. Lipinski's Rule factors multiple qualities of a compound, such as molecular mass and hydrogen bond donors, to determine if it is permeable. Following the maximum of 5 hydrogen bond donors, the maximum of hydrogen bond acceptors, the maximum LogP value of 5, and the maximum molecular mass of 500 daltons, 8 of the 15 compounds are realistically able to be absorbed by the body and circulated to reach proteins, while 7 of the 15 compounds fail to be realistic candidates (Table 5).

Table 5: Qualities of ChemDiv compounds and passing/failing of Lipinski's Rule. 8 candidates "pass" and 7 candidates "fail".

ChemDiv	Hydrogen Bond Donors	Hydrogen Bond Acceptors	Molecular Mass	Consensus LogP	ESOL Water Solubility	Lipinski's Rule
4393-0009	2	7	429.41	3.41	-4.56	Yes
V029-1633	1	6	414.84	3.50	-5.34	Yes
8017-4590	1	6	441.51	2.03	-4.15	Yes
V015-7427	1	6	556.61	4.20	-6.41	No
V027-6087	1	6	392.41	3.08	-4.84	Yes
V027-2157	1	7	398.39	3.02	-4.63	Yes
V015-2773	1	5	532.65	4.30	-6.36	No
K305-0525	1	7	542.60	2.11	-3.81	No
V013-5298	1	5	532.65	4.30	-6.36	No
C260-0238	1	6	475.49	2.81	-4.62	Yes
G923-0287	2	8	530.52	4.92	-6.22	No
V010-4661	1	5	534.67	4.33	-6.53	No
V004-4933	1	7	608.66	4.15	-6.59	No
L299-0508	2	5	466.55	3.31	-4.91	Yes
L299-3577	2	7	460.55	2.41	-4.11	Yes

DOI: 10.36838/IJHSR820.5

Toxicity:

As a drug cannot be fully "safe", the best method to measure hazard is the level of toxicity, and whether the human body can handle the compound without experiencing serious side effects. Thus, the LD50 dosage observes how much dosage of the compound would be lethal for half of a test sample, meaning that a higher LD50 indicates a safer drug. The predicted toxicity class describes a range of values, with the class of 4 being mildly hazardous but not acutely dangerous. That said, the generally accepted LD50 dosage is 500 mg/kg, but anything above 1000 is considered more ideal.[40] While all 8 compounds are still viable, the compounds ChemDiv-V027-2157, ChemDiv-L299-0508, and ChemDiv-L299-3577 have LD50 scores close to the more hazardous 500 mg/kg dosage and so cannot be considered successful candidates (Table 6). Of the 8 viable candidates, ChemDiv-V027-6087 is the most ideal, as the LD50 dosage is safely at 1000 mg/kg, and the molecular dock-

Table 6: Level and kinds of toxicity of each compound, followed by a figure comparing each drug to FDA-approved toxicity levels. The most ideal candidate is ChemDiv-V027-6087.

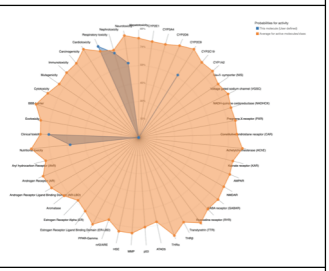
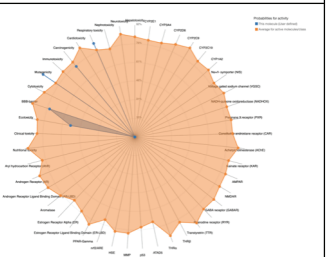
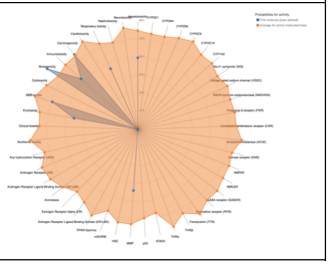
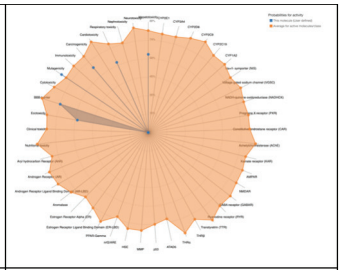
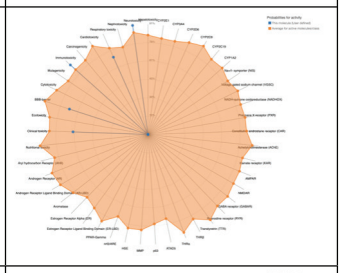
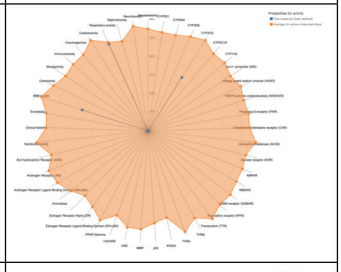
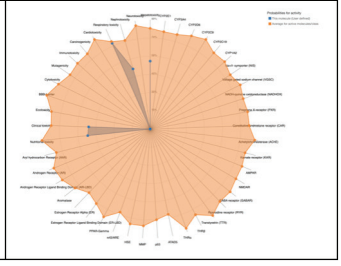
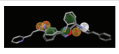
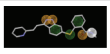
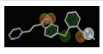
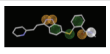
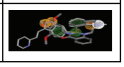
ChemDiv	Predicted LD50 (mg/kg)	Predicted Toxicity Class	Figure
4393-0009	1230	4	
V029-1633	1000	4	
8017-4590	1000	4	
V027-6087	1000	4	
V027-2157	649	4	
C260-0238	2000	4	
L299-0508	560	4	
L299-3577	630	4	

Table 7: Structures of the 5 small molecule drug candidates.

ChemDiv	V027-6087	4393-0009	V029-1633	8017-4590	C260-0238
Structure					

Conclusion

The results of this project were successful in identifying 5 small-molecule drug candidates to target VEGF as potential treatments for vascular malformations. The overall ideal compound is ChemDiv-V027-6087, as its results in each of the experiments were highest when taken from a holistic point of view. With the best molecular docking SwissParam score of -10.0058 and a safe 1000 LD50 dosage, the candidate is able to bind, permeate, and inhibit the VEGF protein to reduce vascular malformation symptoms and stagnate further development. These results provide a new foundation for further research in long-term therapeutics for those with vascular malformations. However, it must be noted that as the majority of these experiments were conducted through computer simulations and in silico models, further validation through *in vitro* and *in vivo* studies must be conducted to account for complex biological variables, such as off-target effects, systemic toxicity, and metabolic stability.

References

1. Adams, D. M.; Trenor, C. C., III; Hammill, A. M.; Vinks, A. A.; *et al.* Efficacy and Safety of Sirolimus in the Treatment of Complicated Vascular Anomalies. Pediatrics 2016, 137 (2), e20153257. https://doi.org/10.1542/peds.2015-3257.
2. Tran, A. Q.; Shoji, M. K.; Levitt, A.; Lee, W. W.; *et al.* Vascular Endothelial Growth Factor Receptor Expression in Orbital Cavernous Malformations and Lymphatic Malformations. *Ophthalmology and Vision Care* 2021, *1* (1), 1003.
3. Muster, R.; Ko, N.; Smith, W.; Su, H.; *et al.* Proof-of-Concept Single-Arm Trial of Bevacizumab Therapy for Brain Arteriovenous Malformations. *Stroke* 2021, *52* (3), e123–e127. https://doi.org/10.1136/bmjno-2020-000114.
4. Bates, D. O. Vascular Endothelial Growth Factors and Vascular Permeability. Cardiovasc. Res. 2010, 87 (2), 262–271. https://doi.org/10.1093/cvr/cvq105.
5. Bhari, N.; Agarwal, A.; Asritha, C. V. V.; Panda, M.; *et al.* Vascular Malformations. *Indian Dermatology Online Journal* 2024, *15* (3), 415–430. https://doi.org/10.4103/idoj.idoj_633_23.
6. Walker, E. J.; Su, H.; Shen, F.; Jun, K.; *et al.* Bevacizumab Attenuates VEGF-Induced Angiogenesis and Vascular Malformations in the Adult Mouse Brain. *Stroke* 2012, *43* (7), 1925–1930. https://doi.org/10.1161/STROKEAHA.111.647982.
7. Snellings, D. A.; Gallione, C. J.; Clark, D. S.; Vozoris, N. T.; *et al.* Somatic Mutations in Vascular Malformations of Hereditary Hemorrhagic Telangiectasia Result in Bi-allelic Loss of *ENG* or *ACVRL1*. *Am. J. Hum. Genet*. 2019, *105* (5), 894–906. https://doi.org/10.1016/j.ajhg.2019.09.010
8. Kirkorian, A. Y.; Grossberg, A. L.; Püttgen, K. B.; Pepin, M. G.; *et al.* Genetic Basis for Vascular Anomalies. *Semin. Cutaneous Med. Surg*. 2016, *35* (3), 128–136. https://doi.org/10.12788/j.sder.2016.051.
9. Queisser, A.; Seront, E.; Boon, L. M.; Vikkula, M.; *et al.* Genetic Basis and Therapies for Vascular Anomalies. *Circ. Res*. 2021, *129* (1), 155–173. https://doi.org/10.1161/CIRCRESAHA.121.318145.
10. Seront, E.; Queisser, A.; Boon, L. M.; Vikkula, M.; *et al.* Molecular Landscape and Classification of Vascular Anomalies. *Hematology Am Soc Hematol Educ Program* 2024, *2024* (1), 700–708. https://doi.org/10.1182/hematology.2024000598.
11. Mansur, A.; Radovanovic, I. Vascular Malformations: An Overview of Their Molecular Pathways, Detection of Mutational Profiles and Subsequent Targets for Drug Therapy. *Front. Neurol.* 2023, *14*, 1099328. https://doi.org/10.3389/fneur.2023.1099328.
12. Seebauer, C. T.; Wiens, B.; Hintschich, C. A.; Batista da Silva, N. P.; *et al.* Targeting the Microenvironment in the Treatment of Arteriovenous Malformations. *Angiogenesis* 2023, *27* (1), 91–103. https://doi.org/10.1007/s10456-023-09896-3.
13. Borst, A.; Vinks, A. A.; Ricci, K.; Hammill, A. M.; *et al.* Targeted Medical Therapies for Vascular Anomalies. *Hematology Am Soc Hematol Educ Program* 2024, *2024* (1), 709–717. https://doi.org/10.1182/hematology.2024000599.
14. Hammer, J.; Seront, E.; Duez, S.; Dupont, S.; *et al.* Sirolimus Is Efficacious in Treatment for Extensive and/or Complex Slow-Flow Vascular Malformations: A Monocentric Prospective Phase II Study. *Orphanet J. Rare Dis*. 2018, *13*, 191. https://doi.org/10.1186/s13023-018-0934-z.
15. van der Vleuten, C. J. M.; Kater, A.; Wijnen, M. H. W. A.; Schultze Kool, L. J.; *et al.* Effectiveness of Sclerotherapy, Surgery, and Laser Therapy in Patients with Venous Malformations: A Systematic Review. *Cardiovasc. Intervent. Radiol.* 2014, *37* (4), 977–989. https://doi.org/10.1007/s00270-013-0764-2.
16. Kane, G.; Fernandez-Pineda, I.; *et al.* Targeted Therapies for Vascular Malformations. *Front. Med*. 2024, *11*, 1446046. https://doi.org/10.3389/fmed.2024.1446046
17. Bao, P.; Kodra, A.; Tomic-Canic, M.; Golinko, M. S.; Ehrlich, H. P.; Brem, H.; *et al.* The Role of Vascular Endothelial Growth Factor in Wound Healing. *J. Surg. Res*. 2009, *153* (2), 347–358. https://doi.org/10.1016/j.jss.2008.04.023
18. Zerbib, L.; Ladraa, S.; Fraissenon, A.; Bayard, C.; *et al.* Targeted Therapy for Capillary-Venous Malformations. *Signal Transduct. Target. Ther.* 2024, *9*, 146. https://doi.org/10.1038/s41392-024-01862-9.
19. Delestre, F.; Venot, Q.; Bayard, C.; Fraissenon, A.; *et al.* Alpelisib Administration Reduced Lymphatic Malformations in a Mouse Model and in Patients. *Sci. Transl. Med.* 2021, *13* (614), eabg0809. https://doi.org/10.1126/scitranslmed.abg0809.
20. Kobialka, P.; Sabata, H.; Vilalta, O.; Graupera, M.; *et al.* The Onset of PI3K-Related Vascular Malformations Occurs during Angiogenesis and Is Prevented by the AKT Inhibitor Miransertib. *EMBO Mol. Med.* 2022, *14* (7), e15619. https://doi.org/10.15252/emmm.202115619.
21. Wang, K.; Zhao, S.; Liu, B.; Zhang, Q.; *et al.* Perturbations of BMP/TGF-β and VEGF/VEGFR Signalling Pathways in Non-Syndromic Sporadic Brain Arteriovenous Malformations. *J. Med. Genet.* 2018, *55* (10), 675–684. https://doi.org/10.1136/jmedgenet-2017-105224.
22. Revencu, N.; Boon, L. M.; Mulliken, J. B.; Clapuyt, P.; *et al.* Vascular Anomalies of the Head and Neck: A Review of Genetics. *Semin. Pediatr. Surg.* 2014, *23* (4), 227–232. https://doi.org/10.1053/j.sempedsurg.2014.06.016.
23. Kalailingam, R. P.; Rannikmae, K.; Hausman-Kedem, M.; Musolino, P. L.; *et al.* Genetic Insights into Hemorrhagic Stroke and Vascular Malformations: Pathogenesis and Emerging Therapeutic Strategies. *Stroke* 2025, *56* (5), 1298–1311. https://doi.org/10.1161/STROKEAHA.124.045182
24. Coulie, J.; Seront, E.; Vikkula, M.; Boon, L. M. Extracranial Arteriovenous Malformations: Towards Etiology-Based Therapeutic Management. *J. Clin. Invest.* 2025, *135* (6), e172837. https://doi.org/10.1172/JCI172837
25. Lee, C.; Kim, M.-J.; Kumar, A.; Lee, H.-W.; *et al.* Vascular Endothelial Growth Factor Signaling in Health and Disease: From Molecular Mechanisms to Therapeutic Perspectives. *Signal Transduct. Target. Ther.* 2025, *10* (1), 170. https://doi.org/10.1038/s41392-025-02249-0.
26. Fährrolfes, R.; Bietz, S.; Flachsenberg, F.; Meyder, A.; Nittinger, E.; Otto, T.; Volkamer, A.; Rarey, M. ProteinsPlus: A Web Portal for Structure Analysis of Macromolecules. Nucleic Acids Res. 2017, 45 (W1), W337–W343. https://doi.org/10.1093/nar/gkx333.
27. Ngan, C.-H.; Hall, D. R.; Zerbe, B.; Grove, L. E.; Kozakov, D.; Vajda, S. FTSite: High Accuracy Detection of Ligand Binding Sites on Unbound Protein Structures. Bioinformatics 2012, 28 (2), 286–287. https://doi.org/10.1093/bioinformatics/btr651.
28. Jendele, L.; Krivák, R.; Skoda, P.; Novotný, M.; Hoksza, D. PrankWeb: A Web Server for Ligand Binding Site Prediction and Visualization. Nucleic Acids Res. 2019, 47 (W1), W345–W349. https://doi.org/10.1093/nar/gkz424.
29. Sunseri, J.; Koes, D. R. Pharmit: Interactive Exploration of Chemical Space. Nucleic Acids Res. 2016, 44 (W1), W442–W448. https://doi.org/10.1093/nar/gkw287.
30. Daina, A.; Michielin, O.; Zoete, V. SwissADME: A Free Web Tool to Evaluate Pharmacokinetics, Drug-Likeness and Medicinal Chemistry Friendliness of Small Molecules. Sci. Rep. 2017, 7, 42717. https://doi.org/10.1038/srep42717.
31. Banerjee, P.; Eckert, A. O.; Schrey, A. K.; Preissner, R. ProTox-II: A Webserver for the Prediction of Toxicity of Chemicals. Nucleic Acids Res. 2018, 46 (W1), W257–W263. https://doi.org/10.1093/nar/gky318.

DOI: 10.36838/IJHSR820.5

32. Kozakov, D.; Grove, L. E.; Hall, D. R.; Bohnuud, T.; *et al.* The FT-Map Family of Web Servers for Determining and Characterizing Ligand-Binding Hot Spots of Proteins. *Nat. Protoc.* 2015, *10* (5), 733–755. https://doi.org/10.1038/nprot.2015.043.
33. Hajduk, P. J.; Huth, J. R.; Fesik, S. W. Druggability Indices for Protein Targets Derived from NMR-Based Screening Data. *J. Med. Chem.* 2005, *48* (7), 2518–2525. https://doi.org/10.1021/jm049131r.
34. Hopkins, A. L.; Groom, C. R.; Alex, A. Ligand Efficiency: A Useful Metric for Lead Selection. *Drug Discov. Today* 2004, *9* (10), 430–431. https://doi.org/10.1016/S1359-6446(04)03069-7.
35. Schinkel, A. H.; Wagenaar, E.; van Deemter, L.; Mol, C. A.; Borst, P. Absence of the Mouse mdr1a P-Glycoprotein Leads to a High Oral Bioavailability and Brain Penetration of Many Drugs. *Proc. Natl. Acad. Sci. U.S.A.* 1996, *93* (18), 10346–10351. https://doi.org/10.1172/JCI118699.
36. Chen, Z.; Li, H.-L.; Zhang, Q.-J.; Bao, X.-G.; *et al.* Pharmacophore-Based Virtual Screening versus Docking-Based Virtual Screening: A Benchmark Comparison against Eight Targets. *Acta Pharmacol. Sin*. 2009, *30* (12), 1694–1708. https://doi.org/10.1038/aps.2009.159
37. Sadremomtaz, A.; Ali, A. M.; Jouyandeh, F.; Balalaie, S.; *et al.* Molecular Docking, Synthesis, and Biological Evaluation of a Vascular Endothelial Growth Factor (VEGF) B-Based Peptide as an Antiangiogenic Agent Targeting the Second Domain of the Vascular Endothelial Growth Factor Receptor 1 (VEGFR1D2) for Anticancer Application. *Signal Transduct. Target. Ther.* 2020, *5*, 76. https://doi.org/10.1038/s41392-020-0177-z.
38. Ramírez D, Caballero J. Is It Reliable to Take the Molecular Docking Top Scoring Position as the Best Solution without Considering Available Structural Data?" *Molecules*. 2018, *23* (5), 1038. https://doi.org/10.3390/molecules23051038
39. Benet, L. Z.; Hosey, C. M.; Ursu, O.; Oprea, T. I. BDDCS, the Rule of 5 and Drugability. *Adv. Drug Deliv. Rev.* 2016, *101*, 89–98. https://doi.org/10.1016/j.addr.2016.05.007.
40. Wang, Y.; Cui, C.; Ren, X.; Dong, X.; Cui, W.; *et al.* Cardiovascular Toxicity Associated with Angiogenesis Inhibitors: A Comprehensive Pharmacovigilance Analysis Based on the FDA Adverse Event Reporting System Database from 2014 to 2021. *Front. Cardiovasc. Med*. 2022, *9*, 988013. https://doi.org/10.3389/fcvm.2022.988013.

Author

Andrew Wi is a senior at Archbishop Mitty High School with a passion for chemistry and its biomedical applications, such as drug discovery. His work targets the vascular endothelial growth factor (VEGF) for the treatment of vascular malformations by using methods such as virtual screening and molecular docking to find a small molecule that can inhibit the VEGF protein. He plans to continue his education in biochemistry and medicine to discover alternative, better treatments for similar diseases. In his spare time, Andrew serves as a representative on his school's student government, captain of his school's speech and debate, and plays clarinet and tennis for his school's Wind Ensemble and Men's Varsity Tennis Team.

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

The Effect of PIEB vs CEI for Quality of Parturient Care and Neonatal Outcomes: A Comparative Review

Angelica C. Cheng

Dulwich College, 71 Bukit Batok West Avenue 8, 658966 Singapore; Angelicacheng814@icloud.com

ABSTRACT: Epidurals are the gold standard for pain relief in labor analgesia worldwide. However, reports on the effects of PIEB compared to CEI on rescue bolus, motor block, and adverse effects remain unclear. This research paper aims to review the effect of PIEB and CEI on rescue bolus, motor block, and adverse effects. Existing literature found that there was an overall lower incidence of requirement for rescue bolus, similar levels of motor block, and either lower or no statistically significant difference in PIEB compared to CEI in terms of adverse effects. Additionally, higher maternal satisfaction was reported in PIEB compared to CEI. Overall, more research is needed to determine whether the results from the studies reveal a general trend or if one or two studies are outliers. Further research on the varying effects of PIEB and CEI can help healthcare suppliers improve pain relief for patients as well as long-term neonatal outcomes.

KEYWORDS: Health and Medicine, Labor Analgesia, Epidural, Programmed Intermittent Epidural Bolus (PIEB), Continuous Epidural Infusion (CEI).

Introduction

Approximately 135 million births occur annually worldwide.[1] Global epidural rates vary, but are usually lower in developing countries.[2] For example, China has an epidural rate of 30%,[3] whereas for some developed countries like Finland, the rate can be as high as 83.4%.[4] Continuous epidural infusion (CEI) is the traditional epidural method of use, where anesthesia is continuously dripped through a catheter inserted into the epidural space in the spine.[5] CEI is associated with increased risk of hypertension, motor block, and numerous other complications, yet it remains the industry's standard.[5] Perhaps this may contribute to the rapidly declining global fertility rate, with the drastic reduction of 5.31-2.25 live births/woman globally from 1963 to 2023.[6] Thus, there has recently been a rise in research related to new epidural methods that aim to improve the quality of pain relief and minimize side effects.[5]

Programmed intermittent epidural bolus (PIEB) is a modern epidural method in which a fixed amount of anesthetic is administered to the patient at pre-programmed intervals using a computer-controlled pump.[7] PIEB has been hypothesized to reduce rescue bolus (amount of analgesia required to reach pain relief), motor block (interruption of motor nerve fibers which allows for pain relief),[8] and adverse effects (side effects incurred due to the analgesia) more than CEI.[9] However, the exact effect on rescue bolus, motor block, and adverse effects between PIEB and CEI remains unclear.

Therefore, the primary purpose of this review is to compare the effects of PIEB versus CEI on rescue bolus, motor block, and adverse effects. The secondary aim of this review is to understand the effect of PIEB versus CEI on women of different parities across parities and maternal satisfaction. By thoroughly reviewing the existing literature surrounding modern epidural methods, this further educates providers on the effect of PIEB across parities, subsequently allowing patient anesthesia to be better tailored to their unique needs and effectively relieving pain during labor and delivery.

Methods

Literature Search Strategy:

PubMed and Google Scholar were the databases used to search for existing literature for review. These databases were utilized from the 5th of October to the 11th of December 2025. Search terms used included "labor," "analgesia," "PIEB," "CEI," "rescue bolus," "motor block," "adverse effects," "parity," "compare," "multicenter," "randomized controlled trial," "quasi-experimental," and "maternal satisfaction."

After applying the search terms mentioned above, 38 studies were determined to be sufficient but required further review. Before screening, 3 duplicate research articles were removed. During the screening process, 28 papers were excluded due to the eligibility criteria. After eligibility criteria were met, 7 papers were included in the review (Figure 1).

Eligibility Criteria:

Randomized and quasi-experimental studies that included patients from obstetric populations who requested labor analgesia and agreed to take part in studies that compared continuous epidural infusions (CEI) and programmed intermittent epidural bolus (PIEB) were included in this review.

Additionally, eligible studies included women between 18 and 40 years of age, ≥37 weeks' gestation, with singleton pregnancies and cephalic presentation, without contradictions to labor analgesia. Studies included women with fetuses that had a healthy baseline fetal status (normal fetal heart rate, absence of decelerations, normal amniotic fluid, and good fetal movement). Lastly, only studies that directly compared the effect of PIEB and CEI on rescue bolus, motor block, and adverse effects were included.

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.6

Non-eligible studies included women with multiple gestation, preterm labor, non-cephalic presentation, unhealthy fetal baseline status, and contradictions to labor analgesia. Additionally, studies that did not directly compare the effect of PIEB and CEI on the effect of rescue bolus, motor block, and adverse effects were excluded.

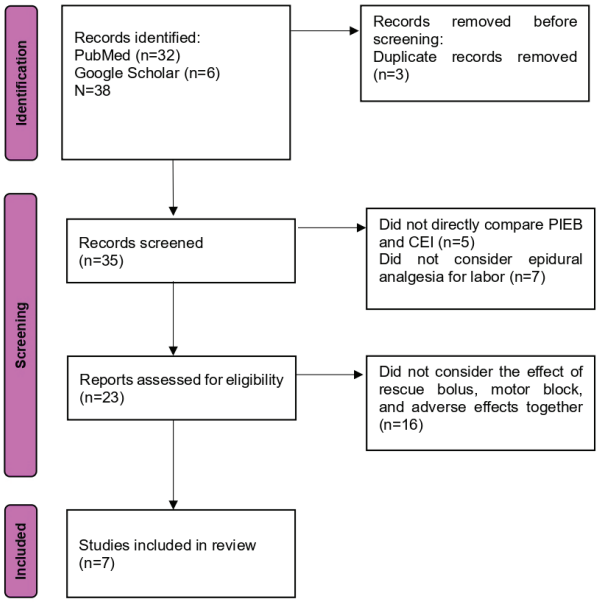


Figure 1: Review inclusion and exclusion criteria PRISMA flowchart. 38 records identified, 35 screened, 23 assessed for eligibility, and 7 studied were included in the review after screening.

Discussion

Rescue Bolus:

Dangers of Excessive Rescue Bolus:

Rescue bolus is the additional amount of labor analgesia required for pain relief to occur. Lower rescue bolus requirement is important for labor analgesia as it means pain relief can be achieved with lower analgesic consumption. This is cost-effective and minimizes potential adverse effects from the rescue boluses. For example, if high volumes of epidural analgesia are administered, it can lead to severe consequences such as stroke due to increased epidural pressure (ICP).[10] This increase in ICP may rupture epidural blood vessels and nerve tissue, which could result in ischemia (reduced blood flow) of the cauda equina (nerve roots extending from the lower end of the spinal cord)[11] or distal spinal cord (lower spinal cord),[12] resulting in neurological deficits (CNS impairment) in patients.[13]

Requirement for Rescue Bolus:

Of the seven studies included in this review, six reported that PIEB reduced rescue bolus requirement more than CEI; however, whether this was statistically significant varied between studies.[9,14-18] One study reported that CEI reduced rescue bolus requirement more compared to PIEB.[19] For example, in 2011, Capogna and colleagues compared the effect of PIEB and CEI on labor analgesia in terms of incidence of motor block and labor outcomes.[9] These researchers found that PIEB had a statistically significantly lower number of patient-controlled epidural analgesia (PCEA) boluses per patient and number of patients requiring additional PCEA boluses ($p<0.001$) compared to CEI.[9] Additionally, in 2024 Ou and Chen investigated the amount of work in obstetrics anesthesia units between PIEB and CEI and found that PIEB had a statistically significantly lower proportion of patients (females about to give birth) that required unscheduled visits from healthcare professionals ((20.8% to 27.7%) ($p=0.033$) (RR=0.69, 95% CI=(0.48-0.97)) compared to CEI (Table 1).[15] Unscheduled visits are defined as "unplanned visits by healthcare staff to control issues outside routine scheduled checks" e.g., requirement for rescue bolus and side effects.[15] Additionally, Haidl *et al.*, in 2022, investigated the effect of PIEB compared to CEI on labor analgesia using an adrenaline epidural solution and reported that the mean number of PCEA boluses required per patient was statistically significantly lower for PIEB compared to CEI (3.9 (SD 4.1) vs 1.9 (SD 2.0)) ($p<0.001$).[14] Lin *et al.*, in 2016, also investigated differences in PIEB and CEI as it relates to the safety of techniques for patients and neonates; the authors reported a statistically significantly lower requirement of PCEA rescue bolus in PIEB compared to CEI (1.27±0.26 vs 5.71±1.03) ($p=0.00$).[17] Lastly, Liu *et al.*, in 2020, determined the efficacy and safety of PIEB and CEI and found there was a statistically significant decrease in total rescue bolus requirement from PIEB compared to CEI ($p=0.030$) (RR=0.61, 95% CI).[20] These studies all suggest a trend that the requirement for rescue bolus was statistically significantly lower in PIEB groups compared to CEI. However, a 2024 study by Kim *et al.* compared PIEB and CEI on the effects of the flow rates for labor analgesia and found that PIEB did not have a statistically significantly lower percentage of patients requiring rescue bolus (65.5% in PIEB vs 78.9% in CEI) ($p=0.081$) compared to CEI.[16] Furthermore, Wang *et al.*, in 2021, which compared PIEB and CEI for duration of labor, mode of delivery, side effects, and maternal satisfaction score, found a statistically significant difference in the number of patients who required PCEA bolus, with PIEB being higher than CEI (425 in PIEB vs 423 in CEI) (OR = 0.27, 95% CI=(0.14–0.51), p <0 .0001, I2 = 65%, p for heterogeneity =0 .01).19 This is the converse to the other studies, as in this case, PIEB required more rescue bolus compared to CEI. This key difference in results in the Wang *et al.* study[19] may be because this study includes several variables that are not accounted for in the other studies and thus, evaluates the relationship between the amount of rescue bolus required and pain relief more thoroughly. Therefore, further research is required to understand whether this study was an anomaly of the trend or whether other studies do not have enough variables included in the analysis to detect accurate differences if they do exist. Overall, these findings show a consensus of reduced rescue bolus requirement with the use of PIEB compared to CEI.

Motor Block:

Impact of Motor Block on the Labor Experience:

Motor block occurs when labor analgesia temporarily blocks nerve signals to muscles, causing reduced or temporary immobility in the legs.[21] This block targets specific spinal levels (e.g., T10-lower sacral nerve roots (S2-4)),[21] which reduces

the patient's ability to push during labor as contractions of the abdominal and pelvic floor muscles are weakened, which could prolong the secondary stage of labor.[22] Such weakened muscle contraction can be harmful to both the patient (more exhaustion and increased chance of instrumental delivery) and fetus (lower Apgar scores or shoulder dystocia).[23] Lower incidence of motor block is extremely important to prioritize due to its ability to improve the quality of the parturient's birth. By reducing motor block, this allows the patient to have more mobility and to find a comfortable birthing position, as well as being able to actively push out the fetus during the second stage of labor.[24]

Incidence of Motor Block:

Of the seven papers referenced above, two found that PIEB had a statistically significantly lower incidence of motor block compared to CEI,[15] one reported PIEB had a statistically significantly higher incidence of motor block compared to CEI,[19] while four reported no statistically significant difference between the two methods of epidural analgesia for labor.[14-16,17-20] For example, Capogna *et al.* found that PIEB had a statistically significantly lower incidence of motor block compared to CEI (2.7% vs 37%) (p<0.001) (OR=21.2, 95% CI=(4.9-129.3), HR=7.8, 95% CI=(1.9-30.8)) (Table 1).[9] Additionally, Ou and Chen found there to be a statistically significantly lower incidence of motor block in PIEB compared to CEI (7.1% vs 2.1%) (p=0.03).[15] Together, these results suggest that PIEB is more favorable compared to CEI in terms of minimizing the incidence of motor block. Conversely, Wang *et al.* found that PIEB was more prone to extensive motor block than CEI.[19] However, Haidl *et al.* found that there was no statistically significant difference in motor block for PIEB compared to CEI at delivery.[14] Similarly, Kim *et al.* reported there was no statistically significant difference in motor blockade between PIEB and CEI after 4 hours of labor analgesia (p=0.224).[16] Furthermore, Lin *et al.* reported no difference in motor block for PIEB compared to CEI.[17] Lastly, Liu *et al.* found that PIEB had a lower risk of motor blockade than CEI; however, these findings were not statistically significant. Thus, there was no statistically significant difference between PIEB and CEI for the risk of motor blockade (p=0.058) (95% CI = (0.19–1.03, RR=0.44) (Table 1).[20] Despite not being statistically significant, there was a 56% reduction in motor blockade, which may be due to over-adjustment of covariates, which may have adjusted away a possible true difference. Over-adjustment occurs when mediators are included, which results in biased results, usually towards the null. Therefore, the results of these 4 studies suggest that there is no statistically significant difference between PIEB and CEI usage in terms of the incidence of motor block. Overall, due to the varying reports of these studies, a larger, more in-depth multicenter study is needed to determine whether PIEB is associated with a higher, lower, or no statistically significant difference in motor block incidence than CEI.

Table 1: Characteristics of research articles included in the review.

Author(s)	Year of Publication	Country and Facility	Research Objective(s)	Key Findings
Capogna et al.[9]	2011	Rome, Italy, Città di Roma Hospital	Compare the incidence of motor block and labor outcome in women who received PIEB or CEI for maintenance of labor analgesia.	**Rescue bolus:** PIEB had a lower number of PCEA boluses per patient and number of patients requiring additional PCEA boluses than CEI (p<0.001). **Motor block:** PIEB had lower motor block compared to CEI (2.7% vs 37%) (p<0.001) (OR=21.2, 95% CI=(4.9-129.3), HR=7.8, 95% CI=1.9-30.8)). **Adverse effects:** None reported.
Haidl et al.[14]	2022	Norway, Akershus University Hospital	Investigate the effect of programmed intermittent boluses using an epidural solution containing adrenaline.	**Rescue bolus:** Mean number of PCEA boluses required per patient was lower for PIEB compared to CEI (3.9±4.0 vs 1.9±2.0) (p<0.001). **Motor block:** No difference in motor block for PIEB compared to CEI at delivery. **Adverse effects:** No significant difference in need for anesthetist intervention. **Maternal satisfaction:** No statistically significant difference in maternal satisfaction between PIEB compared to CEI.
Ou and Chen[15]	2024	Taipei, Taiwan, Shin Kong Memorial Wu Ho-Su Hospital	Investigate the validity that PIEB may function as a feasible tool in reducing the amount of work in the obstetrics anesthesia units.	**Rescue bolus:** PIEB reduced proportion of parturients requiring unscheduled visits* compared to CEI from 27.7% to 20.8% (p=0.033) (OR=0.69, 95% CI=(0.48-0.97)). PIEB had a significant predictive value of maintenance for the probability of requiring unscheduled visits compared to CEI (p<0.01) (OR=0.53, 95% CI=(0.36–0.80). **Motor block:** Lower incidence of motor block in PIEB compared to CEI (7.1% vs 12.1%) (p=0.03). **Adverse effects:** PIEB had a lower rate of episiotomy than CEI regardless of parity, 210 (92.5%) (p<0.01) compared to 159 (83.2%). Degree of perineal laceration lower in PIEB 2(2-3) for nulliparous women compared to CEI 3(2-3) (p<0.01), no difference for parous women. No difference in Apgar score at 1 and 5 mins for both PIEB and CEI for both parous and nulliparous women. **Maternal satisfaction:** Lmaternal satisfaction in PIEB compared to CEI (p=0.0003) (WMD=0.91).
Wang et al.[19]	2021	Multiple Countries (systematic review and meta-analysis)	Compare the efficacy of PIEB with that of CEI with respect to the maintenance of epidural analgesia during labor, duration of labor, mode of delivery, side effects, and satisfaction score.	**Rescue bolus:** Different number of patients who require PCEA bolus, 425 in PIEB compared to 423 in CEI (OR=0.27, 95% CI=(0.14–0.51), p<0.0001, I^2=65%, p for heterogeneity=0.01). **Motor block:** PIEB was more prone to extensive motor block. **Adverse effects:** No difference in nausea, vomiting, need for norepinephrine therapy, and time to first. Lower Apgar score at 1 and 5 minutes for CEI compared to PIEB (WMD=0.07, 95% CI: 0.02-0.13, p=0.007, I2=55%, p for heterogeneity=0.04), (WMD=−0.08, 95% CI=(−0.12–−0.05, p<0.00001, I2=21%, p for heterogeneity=0.27). Maternal satisfaction: Higher maternal satisfaction in PIEB compared to CEI (WMD=0.91, 95% CI=(0.42–1.39), p=0.0003, I^2=98%, p for heterogeneity<0.00001).
Kim et al.[16]	2024	South Korea, Samsung Medical Centre	Compare the effects of epidural analgesics infused at different flow rates on labor analgesia	**Rescue bolus:** No difference in percentage of patients requiring rescue bolus between PIEB and CEI. **Motor block:** No difference in motor blockade 4h after labor analgesia induction. **Adverse effects:** Lower instrumental rate in PIEB compared to CEI (0 out of 29 vs 3 out of 28). 3 participants (2 in CEI and 1 in PIEB) complained of numbness in the lower extremities immediately after labor analgesia induction, they recovered within 4hrs. No difference in Apgar scores at 1 and 5 mins for PIEB compared to CEI (WMD=0.07, 95% CI: 0.02-0.13, p=0.007, I^2=55%, p for heterogeneity=0.04), (WMD=−0.08, 95% CI=(−0.12–−0.05, p<0.00001, I^2=21%, p for heterogeneity=0.27). **Maternal satisfaction:** No statistically significant difference in maternal satisfaction score between PIEB and CEI (9(8,10) in CEI compared to 9(9,10) in PIEB) (p=0.287).
Lin et al.[17]	2016	China, Guangxi Medical University	Investigate whether there are any differences between CEI and programmed IEB analgesia for the Chinese parturients undergoing spontaneous delivery and to further investigate the safety of these techniques for parturients and neonates.	**Rescue bolus:** Lower requirement of PCEA rescue bolus in PIEB compared to CEI (1.27±0.26 vs 5.71±1.03 (p=0.00)). **Motor block:** No difference in motor block for PIEB and CEI. **Adverse effects:** Longer first and second stage labor in CEI compared to PIEB. No difference in labor instruments usage when compared with normal age, Apgar score at 1 minute in PIEB compared to CEI, and Apgar scores at 5 minutes was lower in PIEB compared to CEI but not statistically significant.
Liu et al.[20]	2020	Multiple Countries (meta-analysis)	Determine the efficacy and safety of PIEB and CEI for labor analgesia.	**Rescue bolus:** Decrease in total rescue bolus requirement from PIEB compared to CEI (p=0.030) (RR=0.61). **Motor block:** No difference in PIEB compared to CEI. **Adverse effects:** No difference in pruritus, nausea and vomiting, hypotension, shivering, fetal bradycardia, postpartum hemorrhage, intrapartum fever, perineal tears, and retained placental products between PIEB and CEI. **Maternal satisfaction:** Higher maternal satisfaction using PIEB compared to CEI.

Adverse Effects:

Risk of Adverse Effects:

Adverse effects are unintended side effects that negatively harm the physical or mental health of the parturient or infant.[25] These are important to consider as it allows healthcare professionals to conduct a risk-benefit analysis for the parturient and infant to consider whether the potential complications of using either epidural method outweigh the improved pain relief caused by it.[26]

Impact on Adverse Effects:

Of the seven studies included, three found statistically significant differences in PIEB compared to CEI for adverse effects,[15-19] while four found no statistically significant difference.[9,14,17,20] Firstly, in 2024, Ou and Chen discovered that PIEB had a higher rate of episiotomy than CEI, regardless of parity, 210 (92.5%) compared to 159 (83.2%) (p<0.01).[15] This suggests that PIEB is worse compared to CEI in terms of the occurrence of the adverse effect of an episiotomy. Conversely, the degree of perineal laceration was statistically significantly lower in PIEB 2(2-3) compared to CEI 3(2-3) (p<0.01) for nulliparous women but not statistically significantly the same for parous women, 2(2-2) (p=0.67).[15] Together, these results indicate that PIEB results in fewer adverse effects in mothers compared to CEI after labor. However, in 2021, Wang *et al.* showed that there was no statistically significant difference in nausea, vomiting, need for norepinephrine therapy, and time to first passage of flatus or feces (p=0.06) when comparing PIEB to CEI.[19]

DOI: 10.36838/IJHSR820.6

Additional adverse effects evaluated include numbness, duration of labor, and rate of instrumental delivery. For example, in 2024, Kim *et al.* reported a lower instrumental rate in PIEB compared to CEI (0 out of 29 vs 3 out of 28).[16] Within this same study, of the 85 participants (PIEB=29, CEI=28, and manual=28), 3 participants (1 in PIEB and 2 in CEI) complained of numbness in the lower extremities immediately after labor analgesia induction, but all recovered within 4 hours.[16] In 2011, Capogna *et al.* reported no difference in duration of labor between PIEB and CEI.[9] Similarly, in 2022, Haidl *et al.* found there was no significant difference in need for anesthetist intervention between PIEB and CEI.[14] Additionally, in 2016, Lin *et al.* found there to be no statistically significant difference in labor instrument usage when compared with normal age between the two epidural groups.[17] Furthermore, there was not a statistically significant shorter first and second stage labor in PIEB compared to CEI (420.35±20.29 vs 431.51±19.09) (p=0.75) and (55.31±9.71 vs) (p=0.68).[17] Lastly, in 2020, Liu *et al.* discovered no statistically significant difference in pruritus (p=0.703, RR=1.03), nausea and vomiting (p=0.435, RR=1.44), hypotension (p=0.084, RR 1.43), shivering (p=0.913, RR=0.98), fetal bradycardia (p=0.944, RR 1.04), postpartum hemorrhage (p=0.607, RR=0.73), and intrapartum fever, perineal tears, retained placental products between PIEB and CEI.[20] Overall, these studies do not support the theory that PIEB is more favorable than CEI in terms of maternal adverse effects.

In addition to maternal adverse effects, epidurals can also result in neonatal complications. For example, in 2021, Wang *et al.* outlined how there were statistically significantly higher Apgar scores at 1 and 5 minutes for PIEB compared to CEI (WMD = 0.07, 95% CI: 0.02-0.13, p=0.007, I2 = 55%, p for heterogeneity = 0.04), (WMD = −0.08, 95% CI=(−0.12-−0.05, p < 0.00001, I2 = 21%, p for heterogeneity =0 .27).[19] Additionally, in Ou and Chen, both PIEB and CEI had the same non-statistically significant Apgar score at 1 and 5 minutes for both parous and nulliparous women. At 1 minute, nulliparous women had non-statistically significant Apgar scores of 8(8-9) (p=0.82), while multiparous women had average Apgar scores of 9(8-9) (p=0.36).[15] Similarly, at 5 minutes, there was a nonstatistically significant average Apgar score of 9(8-9) for both nulliparous (p=0.30) and multiparous women (p=0.49).[15] Lastly, in 2024, Kim *et al.* found the same non-statistically significant Apgar scores at 1 minute (PIEB=9(9,9), CEI=9(9,9), manual=9(9,9), p=0.666) and at 5 minutes (PIEB=10(9,10), CEI=10(10,10), manual=10(9,10), p=0.798).[16] Also, there was no statistically significant difference in Apgar score at 1 minute in PIEB compared to CEI (8.62±0.29 vs 8.57±0.16) (p=0.74).[17] Moreover, Apgar scores at 5 minutes were lower in PIEB compared to CEI (9.03±0.18 vs 9.13±0.16) (p=0.06).[17] Together, these studies generally indicate that neither PIEB nor CEI is more favorable when considering the effect they have on neonatal Apgar scores. However, more research is needed to determine if this is true, given that the study did find there to be a statistically significant difference.

Overall, most of the studies indicated there was not a statistically significant difference in adverse effects in both groups, suggesting either can be used when considering side effects for patients and the fetus. However, it is important to consider that most studies considered varying adverse effects, meaning it is difficult to directly compare the adverse effects in each study.

Further Explorations of Maternal and Fetal/Neonatal Adverse Effects of PIEB and CEI:

For this paper, the reviewed literature identified a larger number of maternal effects compared to fetal/neonatal (Table 2). However, the adverse effects on the fetus are critical as it can be directly impacted by the adverse effects of the mother.[27]

An episiotomy is carried out to prevent perineal laceration.[28] Ou and Chen found that although there was a higher episiotomy rate using PIEB, there was ultimately lower severity of perineal lacerations amongst nulliparous women (Table 2).[15] This suggests there are mixed effects on perineal trauma overall. Maternal discomfort can include nausea, vomiting, and itching, and can cause discomfort to the parturient and affect their birthing experience. There was no statistically significant difference in maternal discomfort, such as nausea, vomiting, and itching, in PIEB in both Wang *et al.* and Liu *et al.*[19,20] This indicated similar rates of maternal discomfort in both. Norepinephrine therapy utilizes synthetic norepinephrine to maintain blood pressure in patients suffering from severe hypotension.[29] Wang *et al.* and Heidl *et al.* found no statistically significant difference between PIEB and CEI for norepinephrine therapy and anesthetist intervention, suggesting similar rates of clinically significant hypotension requiring intervention.[14,19] Moreover, heightened duration of labor increases the likelihood of instrumental delivery, infections, and maternal exhaustion.[30] Wang *et al.*, Capogna *et al.*, and Lin *et al.* found no statistically significant difference between PIEB and CEI for time to first passage of feces and duration of first and second stage labor.[9,17,19] Therefore, it can be concluded that there was a similar recovery of gastrointestinal function postpartum in both, and risks associated with longer labor in PIEB and CEI. Instrumental-assisted delivery can cause maternal pelvic injury, postpartum hemorrhage, and infection.[31] Interestingly, while Kim *et al.* reported lower instrumental delivery rates with PIEB compared to CEI, Lin *et al.* found no statistically significant difference between PIEB and CEI.[16,17] This may require further research to conclude whether PIEB has statistically significant benefits in relation to assisted delivery compared to CEI. Excessive postpartum hemorrhage increases blood loss, which results in hypotension.[32] Liu *et al.* and Wang *et al.* both found that there was a statistically significant difference between PIEB and CEI.[18,19] Thus, this indicates similar risks of severe maternal blood loss. High intrapartum fever can increase postpartum hemorrhage and, in severe cases, be fatal, even leading to organ failure.[33] Liu *et al.* found no significant difference between PIEB and CEI for intrapartum fever.[18] In conclusion, while PIEB was associated with higher episiotomy but lower severity of perineal laceration and assisted deliveries, other maternal effects had no statistically significant difference between PIEB and CEI.

In terms of neonatal/fetal effects, fetal bradycardia refers to a sustained fetal heart rate below 110 bpm,which can lead

to hypoxic-ischemic injury and even severe acute hypoxia.[34] Liu *et al.* found no statistically significant difference between PIEB and CEI for fetal bradycardia, suggesting similar fetal cardiovascular safety profiles.[18] Lastly, Apgar scores are a test performed on a neonate 1 and 5 minutes after birth. The 1-minute score determines how well the baby has tolerated the birthing process, while the 5-minute score exemplifies how the neonate is adapting to living outside the mother's womb.[35] Ou and Chen, Kim *et al.*, and Lin *et al.* found no statistically significant difference in Apgar scores at 1 and 5 minutes; however, Wang *et al.* reported slightly higher scores with PIEB compared to CEI.[15-17,19] Overall, Apgar scores are similar in both PIEB and CEI. Conclusively, there is no large statistically significant difference in the neonatal/fetal effects of PIEB and CEI.

Table 2: Maternal and Fetal/Neonatal Adverse Effects of PIEB vs CEI.

	Adverse Effects	Studies	Definitive Conclusion
Maternal	Episiotomy/perineal laceration	Ou and Chen[15]	PIEB was associated with higher episiotomy rate than CEI, but lower severity of perineal lacerations amongst nulliparous women.
	Nausea/vomiting/itching	Wang et al., Liu et al.[19-20]	No statistically significant difference between PIEB and CEI.
	Norepinephrine therapy	Wang et al., Haidl et al.[19-14]	No statistically significant difference between PIEB and CEI.
	Anesthetist intervention		
	Time to first passage of faces	Wang et al., Capogna et al., Lin et al.[19-9-17]	No statistically significant difference between PIEB and CEI.
	Duration of labor		
	Instrument assisted delivery	Kim et al., Lin et al.[16-17]	Kim et al. reported lower instrumental delivery rates with PIEB compared to CEI, Lin et al. found no statistically significant difference between PIEB and CEI.
	Blood pressure	Liu et al., Wang et al.[18-19]	No statistically significant difference between PIEB and CEI.
	Hemorrhage		
	Intrapartum fever	Liu et al.[18]	No statistically significant difference between PIEB and CEI.
Fetal/Neonatal	Fetal bradycardia	Liu et al.[18]	No statistically significant difference between PIEB and CEI.
	Apgar scores	Wang et al., Ou and Chen, Kim et al., Lin et al.[15-16-17-19]	Ou and Chen, Kim et al. and Lin et al. found no statistically significant difference in Apgar scores at 1 and 5 minutes, however Wang et al. reported slightly higher scores with PIEB compared to CEI.

Comparative Effect of PIEB and CEI on Maternal Factors and Outcomes:

Importance of Parity Stratification and Maternal Satisfaction:

Maternal satisfaction scores are a crucial indication of the quality of the labor experience for patients. These scores also allow the parturient to feel more included and that their opinion is valued. Furthermore, having a positive birth experience boosts self-esteem.[36] Conversely, low satisfaction may lead to increased risk of postpartum depression, anxiety, post-traumatic stress disorder (PTSD), and weakened mother-infant bonding.[37] Maternal satisfaction scores are usually collected through questionnaires, for example, the "Childbirth Experience Questionnaire (CEQ)" which contains 22 questions combining a 4-point Likert scale with a visual analog scale (VAS) measured from 0-100 to determine the patient's satisfaction with the labor and delivery process.[38]

Parity refers to the number of pregnancies that reached 20 weeks' gestation regardless of the outcome that a parturient has experienced.[39] Being able to separate maternal and labor outcomes based on parity (parity stratification) is important, as nulliparous women tend to have longer, more challenging labors as their cervix and pelvic tissues are undergoing the labor process, specifically stretching and opening, for the first time.[40] Furthermore, there is a lower risk of obstetric and sphincter injuries (OASIS) in multiparous compared to nulliparous women.[41] Additionally, parous women usually have shorter lengths of labor, so having a mixed parity group and not stratifying by parity may skew conclusions and lead to an underestimation of the length of labor.[41] Lastly, stratifying by parity allows the results of the study to be applied more specifically to all women, which makes the findings applicable to a larger demographic.

Effect on Maternal Outcomes and Parity Consideration:

Of the seven papers included, five papers reported maternal satisfaction,[15-20] while the remaining two did not.[9-17] Firstly, in 2024, Ou and Chen discovered there was a statistically significantly higher maternal satisfaction in PIEB compared to CEI ($p = 0.0003$) (WMD = 0.91).[15] Also, this is the only study that stratified adverse effects by parity.[15] It was found that the degree of perineal laceration was lower in PIEB compared to CEI for nulliparous women but the same for parous women; however, these results were not statistically significant.[15] Furthermore, both PIEB and CEI had the same Apgar score at 1 and 5 minutes for both parous and nulliparous women; however, these results were also not statistically significant.[15] Moreover, in 2021, Wang *et al.* found a statistically significantly higher maternal satisfaction in PIEB compared to CEI (WMD = 0.91, 95% CI = (0.42–1.39), p = 0 .0003, I2 = 98%, p for heterogeneity<0.00001).[19] Lastly, in 2020, Liu *et al.* discovered that using CEI had lower maternal satisfaction compared to PIEB.[20] Together, these three studies indicate that PIEB is more favorable compared to CEI in terms of parturient satisfaction. However, in 2024, Kim *et al.* reported no statistically significant difference in maternal satisfaction score (9(8,10) in CEI compared to 9(9,10) in PIEB) (p=0.287).[16] Additionally, in 2022, Haidl *et al.* found there was no statistically significant difference in maternal satisfaction between PIEB and CEI.[14] Overall, the method to measure maternal satisfaction for each study must be considered, as well as the different preferences of patients regarding their labor experience, as this may explain the varying maternal satisfaction. Similarly, maternal satisfaction is subjective and cannot be quantitatively measured and accurately compared across studies, especially across studies conducted in different facilities and regions that have access to varying qualities ofequipment and training.

■ Strengths and Limitations

This review offers a balanced and unbiased consideration of the compared effects of PIEB and CEI on the incidence of motor block, requirement for rescue bolus, and adverse effects.

DOI: 10.36838/IJHSR820.6

This literature review addresses a key gap in labor analgesia literature. This is because no known literature reviews simultaneously offer information on the requirement for rescue bolus, incidence of motor block, adverse effects, and maternal satisfaction, as well as consider parity stratification. Furthermore, this review included two multicenter studies to improve reliability, as well as studies from multiple countries. This increases the generalizability of the findings to a larger geographical population. Additionally, this review included papers that stratified by parity to allow results to be applied to both parous and nulliparous women.

One of the limitations of this review is the limited number of studies included, this being largely because many studies failed to meet the eligibility criteria. This limits the scope of information that the review can consider, so a strong conclusion cannot be formed based on the limited evidence. Similarly, only one study stratified by parity, a secondary aim of this study, meaning researchers and providers may be wary of trusting the findings based on parity. Lastly, studies included were very heterogeneous and varied in the type of analgesia used, dosage, and point of labor at which it was administered, which makes the results difficult to compare.

Several methods were used to minimize the limitations. Firstly, the eligibility criteria were broadened to allow more research papers to be included. Additionally, to minimize heterogeneity, this review clearly stated the type of analgesia used, dosage, and point of labor at which it was administered. This helps emphasize that the results cannot be directly compared and only suggest a general trend.

Conclusion

Overall, labor analgesia is important to relieve pain during labor and delivery. Two common epidural methods are PIEB and CEI. Of the studies reviewed, 86% reported a rescue bolus, and 67% of the findings were statistically significant in terms of whether PIEB led to a higher or lower requirement for rescue bolus compared to CEI.[9,14,19] There was an overall reduction in rescue bolus in PIEB compared to CEI, apart from one anomaly,[19] which could be due to adjusting covariates. All papers reported the incidence of motor block (43% statistically significant) and adverse effects (43% statistically significant).[15-19] Additionally, 71% reported secondary outcomes (maternal satisfaction or parity stratification), 60% of which were statistically significant.[15-19] Furthermore, motor block was found to be lower in PIEB groups than CEI; however, the statistical significance of these findings varied. Similarly, adverse effects varied between statistically significantly higher in PIEB compared to CEI, as well as no statistically significant difference. Lastly, maternal satisfaction was either statistically significantly higher in PIEB compared to CEI, or there was no statistically significant difference between the two groups. Future studies should stratify by parity and be conducted over multiple facilities, and also choose only a few of the same adverse effects instead of a large range. Furthermore, future studies should avoid including indiscriminate variables by using a theory-driven approach instead to only identify the necessary confounders and minimize overadjustment. Further education on PIEB compared to CEI can improve the quality of parturient care and neonatal outcomes.

Acknowledgments

I am very grateful to my mentor, Bre Calhoun, who gave me advice and helped me develop and edit my research paper. Furthermore, I would like to acknowledge Crimson Education, which connected me to Dr. Calhoun through the Indigo Research program, for whom I am also thankful. Lastly, I want to thank my parents who supported me throughout my research journey.

References

1. Blencowe H, Cousens S, Chou D, *et al.* Born Too Soon: The global epidemiology of 15 million preterm births. *Reprod Health*. 2013; 10(1):S2. doi:10.1186/1742-4755-10-S1-S2
2. Dobson M. Opinion: Labor Analgesia in the Developing World; Why Not. *Rev Pain*. 2011; 5(3):2-3. doi:10.1177/204946371100500302
3. Zhu X, Wang X, Zhang X, Jia P, He G, Chen F. Public attention related to epidural labor analgesia in mainland China: Evidence from internet search data. *Medicine (Baltimore)*. 2025; 104(46):e45978. doi:10.1097/MD.0000000000045978
4. Seijmonsbergen-Schermers AE, van den Akker T, Rydahl E, *et al.* Variations in use of childbirth interventions in 13 high-income countries: A multinational cross-sectional study. *PLoS Med.* 2020; 17(5):e1003103. doi:10.1371/journal.pmed.1003103
5. George RB, Allen TK, Habib AS. Intermittent Epidural Bolus Compared with Continuous Epidural Infusions for Labor Analgesia: A Systematic Review and Meta-Analysis. *Anesth Analg*. 2013; 116(1):133. doi:10.1213/ANE.0b013e3182713b26
6. Roser M. The global decline of the fertility rate. *Our World Data*. Published online February 20, 2014. Accessed December 20, 2025. https://ourworldindata.org/global-decline-fertility-rate
7. Programmed intermittent epidural bolus as an ideal method for labor analgesia: a randomized controlled trial. Accessed November 23, 2025. https://ekja.org/journal/view.php?number=8921
8. Motor Nerve Block - an overview | ScienceDirect Topics. Accessed December 11, 2025. https://www.sciencedirect.com/topics/medicine-and-dentistry/motor-nerve-block
9. Capogna G, Camorcia M, Stirparo S, Farcomeni A. Programmed intermittent epidural bolus versus continuous epidural infusion for labor analgesia: the effects on maternal motor function and labor outcome. A randomized double-blind study in nulliparous women. *Anesth Analg*. 2011; 113(4):826-831. doi:10.1213/ANE.0b013e31822827b8
10. Zabolotskikh I, Trembach N. Safety and efficacy of combined epidural/general anesthesia during major abdominal surgery in patients with increased intracranial pressure: a cohort study. *BMC Anesthesiol.* 2015; 15:76. doi:10.1186/s12871-015-0056-2
11. Kalogeris T, Baines CP, Krenz M, Korthuis RJ. Ischemia/Reperfusion. *Compr Physiol.* 2017; 7(1):113-170. doi:10.1002/j.2040-4603.2017.tb00741.x
12. Cauda equina syndrome - ScienceDirect. Accessed January 11, 2026. https://www.sciencedirect.com/science/article/abs/pii/S0263931924000620
13. Bosscher H. Pressure-Volume Relationships in the Spinal Canal and Potential Neurological Complications After Epidural Fluid Injections. *Front Pain Res.* 2022;3:884277. doi:10.3389/fpain.2022.884277

14. Haidl F, Arne Rosseland L, Rørvik AM, Dahl V. Programmed intermittent boluses vs. continuous epidural infusion in labor using an adrenaline containing solution: A randomized trial. *Acta Anaesthesiol Scand*. 2020; 64(10):1505-1512. doi:10.1111/aas.13689
15. Ou CH, Chen WT. Programmed Intermittent Epidural Bolus Reduces Workloads in Labor Analgesia: A Single Center's Experience. *Medicina (Mex)*. 2024; 60(6):993. doi:10.3390/medicina60060993
16. Kim D, Kim J, Choo H, Choi DH. Programmed intermittent epidural bolus as an ideal method for labor analgesia: a randomized controlled trial. *Korean J Anesthesiol*. 2024; 77(1):106-114. doi:10.4097/kja.23173
17. Lin Y, Li Q, Liu J, Yang R, Liu J. Comparison of continuous epidural infusion and programmed intermittent epidural bolus in labor analgesia. *Ther Clin Risk Manag*. 2016; 12:1107-1112. doi:10.2147/TCRM.S106021
18. Liu SK, Wu SC, Hung SC, Chen KB, Illias AM, Tsai YF. Combined Programmed Intermittent Bolus and Patient-Controlled Bolus Is a More Favorable Setting for Epidural Pain Relief Than Continuous Infusion. *Healthcare*. 2023; 11(9):1350. doi:10.3390/healthcare11091350
19. Wang X xue, Zhang X lan, Zhang Z xia, *et al.* Programmed intermittent epidural bolus in parturients. *Medicine (Baltimore)*. 2022; 101(5):e28742. doi:10.1097/MD.0000000000028742
20. Liu X, Zhang H, Zhang H, Guo M, Gao Y, Du C. Intermittent epidural bolus versus continuous epidural infusions for labor analgesia: A meta-analysis of randomized controlled trials. *PLoS ONE*. 2020; 15(6):e0234353. doi:10.1371/journal.pone.0234353
21. Novikova N, Cluver C. Local anaesthetic nerve block for pain management in labour. *Cochrane Database Syst Rev.* 2012; 2012(4):CD009200. doi:10.1002/14651858.CD009200.pub2
22. Graham AC, McClure JH. Quantitative assessment of motor block in laboring women receiving epidural analgesia. *Anaesthesia*. 2001; 56(5):470-476. doi:10.1046/j.1365-2044.2001.01524-6.x
23. Young C, Bhattacharya S, Woolner A, *et al.* Maternal and perinatal outcomes of prolonged second stage of labor: a historical cohort study of over 51,000 women. *BMC Pregnancy Childbirth*. 2023; 23:467. doi:10.1186/s12884-023-05733-z
24. Antonakou A, Papoutsis D. The Effect of Epidural Analgesia on the Delivery Outcome of Induced Labor: A Retrospective Case Series. *Obstet Gynecol Int*. 2016; 2016:5740534. doi:10.1155/2016/5740534
25. Definition of adverse effect - NCI Dictionary of Cancer Terms - NCI. February 2, 2011. Accessed February 4, 2026. https://www.cancer.gov/publications/dictionaries/cancerterms/def/adverse-effect
26. Liu ZH, Wang DX. Potential impact of epidural labor analgesia on the outcomes of neonates and children. *Chin Med J (Engl)*. 2020; 133(19):2353-2358. doi:10.1097/CM9.0000000000000900
27. Goldenberg RL, Culhane JF, Johnson DC. Maternal Infection and Adverse Fetal and Neonatal Outcomes. *Clin Perinatol*. 2005; 32(3):523-559. doi:10.1016/j.clp.2005.04.006
28. Episiotomy: Procedure, Advantages, Complications & Healing. Cleveland Clinic. Accessed May 10, 2026. https://my.clevelandclinic.org/health/treatments/22904-episiotomy
29. Smith MD, Maani CV. Norepinephrine. In: *StatPearls*. StatPearls Publishing; 2026. Accessed May 10, 2026. http://www.ncbi.nlm.nih.gov/books/NBK537259/
30. Prolonged Labor: Causes, Risks & Treatment. Cleveland Clinic. Accessed May 10, 2026. https://my.clevelandclinic.org/health/diseases/24752-prolonged-labor
31. Forceps or vacuum delivery. nhs.uk. December 2, 2020. Accessed May 10, 2026. https://www.nhs.uk/pregnancy/labour-and-birth/forceps-or-vacuum-delivery/
32. Hemorrhagic Hypotension - an overview | ScienceDirect Topics. Accessed May 10, 2026. https://www.sciencedirect.com/topics/medicine-and-dentistry/hemorrhagic-hypotension
33. Zamstein O, Wainstock T, Sheiner E. Intrapartum Maternal Fever and Long-Term Infectious Morbidity of the Offspring. *J Clin Med.* 2023; 12(9):3329. doi:10.3390/jcm12093329
34. Nayeri JC Unzila. Fetal bradycardia. Cancer Therapy Advisor. January 17, 2019. Accessed May 10, 2026. https://www.cancertherapyadvisor.com/home/decision-support-inmedicine/obstetrics-and-gynecology/fetal-bradycardia/
35. Apgar score: MedlinePlus Medical Encyclopedia. Accessed May 10, 2026. https://medlineplus.gov/ency/article/003402.htm
36. Karlsdottir SI, Leap N. The importance of promoting positive childbirth experiences for women: a perspective paper. *Front Glob Womens Health*. 2025; 6:1599249. doi:10.3389/fgwh.2025.1599249
37. Mukerenge NF, Kafene P, Nakakuwa FN. An assessment of women's satisfaction with maternity care during labor and birth in a referral hospital in Namibia. *Heliyon*. 2025; 11(4):e42644. doi:10.1016/j.heliyon.2025.e42644
38. Patabendige M. Childbirth Experience Questionnaire (CEQ): research proposal for translation and validation into Sinhala language among a Lankan cohort of women. *BMC Res Notes*. 2019; 12:458. doi:10.1186/s13104-019-4499-2
39. Gravidity and Parity - an overview | ScienceDirect Topics. Accessed February 5, 2026. https://www.sciencedirect.com/topics/medicine-and-dentistry/gravidity-and-parity
40. Tilden EL, Snowden JM, Bovbjerg ML, *et al.* The duration of spontaneous active and pushing phases of labor among 75,243 US women when intervention is minimal: A prospective, observational cohort study. eClinicalMedicine. 2022;48:101447. doi:10.1016/j.eclinm.2022.101447
41. Voillequin S, Quibel T, Rozenberg P, Rousseau A. Duration of the second and third stages of labor and risk of postpartum hemorrhage: a cohort study stratified by parity. *BMC Pregnancy Childbirth*. 2025; 25:143. doi:10.1186/s12884-025-07229-4

Author

Angelica is a Year 12 IB student at Dulwich College Singapore and aspires to study medicine in the future. Specifically, she hopes to become an anesthesiologist and specialize in epidurals.

DOI: 10.36838/IJHSR820.6

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Plasma Braking in Subcritical X-ray Pulsars: Coulomb Deceleration in Strongly Magnetized Atmospheres

Chen Xingyu

Raffles Institution, 1 Raffles Institution Lane, Singapore, Singapore, 575954, Singapore; ryanchenxy@gmail.com
Mentor: Alexander Mushtukov

ABSTRACT: We investigate plasma braking in sub-critical X-ray pulsars (XRPs), where radiation pressure is negligible and infalling protons are decelerated primarily by Coulomb collisions in a strongly magnetized neutron-star atmosphere. Our key advance is a self-consistent numerical model that couples the evolution of the flow velocity and the atmospheric density, rather than prescribing a fixed background atmosphere or an ad hoc stopping depth. The model integrates the coupled momentum and deceleration equations while evaluating a generalized Coulomb logarithm applicable across classical, quantum, and strongly magnetized regimes, incorporating the Debye length, quantum diffraction, and electron gyroradius (Larmor radius). We compute stopping lengths and compression profiles across a range of magnetic field strengths, temperatures, densities, and compositions. We find that self-consistent atmospheric compression steepens the upper-layer density gradient and can significantly modify the depth and shape of the braking region compared to fixed-density estimates. These vertical structure profiles have direct implications for radiative transfer and vacuum-resonance mode conversion, providing physically grounded inputs for spectral and polarization modeling of accreting neutron stars.

KEYWORDS: Physics and Astronomy, Astrophysics, Neutron Stars, X-Ray Pulsars, Plasma, Atmospheres.

Introduction

Accretion of matter onto a strongly magnetized neutron star (NS) gives rise to X-ray pulsars (XRPs).[1] In XRPs, the surface magnetic field strength is confirmed to be at least $\gtrsim 10^{12}$ G, which drastically alters both the accretion flow geometry and the microphysics of radiation–matter interaction.[2] The luminosities of XRPs span an enormous range, from about 10^{33} erg s^{-1} up to ~10^{41} erg s^{-1}. At the low end, accretion either halts due to the transition into the "propeller" regime[3-5] or proceeds in a stable fashion from a recombined cold disc.[6] At the highest luminosities, some XRPs are classified as ultra-luminous X-ray sources.[7,8] The strong magnetic field channels matter onto small regions near the magnetic poles of the NS, where most of the gravitational energy is released. The structure of this emission region is determined by the mass accretion rate and corresponding luminosity. At relatively low luminosities ($L \lesssim 10^{37}$ erg s^{-1}), the accretion flow is stopped in the NS atmosphere, and radiation escapes from hot spots of area $\sim 10^9$ cm^2.[9,10] When $L \gtrsim 10^{37}$ erg s^{-1}, the deceleration occurs above the surface in a radiation-dominated shock, producing extended accretion columns.[11-13] Spectra of XRPs are usually modeled with a power law plus a high-energy cutoff. At very low accretion rates ($L \lesssim 10^{35}$ erg s^{-1}), some sources display a spectral transition and develop a two-component spectrum.[14-16] Although models exist for both the low-luminosity regime[17,18] and standard luminosity states,[19] no single theory yet explains the spectral transitions self-consistently. Independent of the model details, the main radiative processes include magnetic Compton scattering,[20] free–free emission and absorption,[21] and cyclotron processes.[22] These processes depend strongly on the polarization state of X-ray photons, especially below the electron cyclotron energy $E_{cyc} \simeq 11.6\ (B/10^{12}\ \mathrm{G})$ keV. Because the opacity is polarization-dependent, high linear polarization degrees (~80%) have been predicted by several XRP models.[18,23-25] However, observations with IXPE26 show systematically lower polarization levels for all XRPs studied so far.[27-33] One factor influencing the emergent polarization is the conversion of polarization modes near the vacuum resonance.[21,34] In a strongly magnetized plasma, radiative transfer is described in terms of two normal modes whose polarization is set by both plasma and vacuum contributions to the dielectric tensor.[35-38] In the dense lower atmosphere, plasma effects dominate, while in the tenuous upper layers, vacuum polarization becomes important. As the modes propagate outward, their polarization vectors can interchange, depending on the local mass density gradient: for shallow gradients, nearly adiabatic O–X or X–O conversion occurs, while steep gradients lead to partial or non-adiabatic mode conversion.[21] A key factor shaping the mass density profile in XRPs is the braking of the accretion flow as it penetrates the atmosphere. The momentum transfer from protons to the atmospheric plasma compresses the upper layers and modifies the density gradient, thereby affecting both radiative transfer and polarization formation. In this work, we focus specifically on the physics of plasma braking in sub-critical XRPs, where the radiative force is negligible, and Coulomb collisions dominate the deceleration process. We investigate how the flow–atmosphere interaction sets the vertical structure and discuss its implications for X-ray emission and polarization. Previous work has shown that Coulomb interactions can efficiently brake the accretion flow in the low-luminosity regime and that the resulting density gradient is important for radiative transfer and polarization formation.

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.7

However, many treatments effectively assume a background atmospheric structure (or prescribe a stopping depth) and then compute deceleration on top of it. In reality, momentum deposition from the decelerating flow compresses the atmosphere, and that compression in turn feeds back on the Coulomb drag through the local density.

The key contribution of this paper is a self-consistent numerical model that couples the evolution of the flow velocity and atmospheric density in sub-critical X-ray pulsars. Unlike previous approaches that prescribe a fixed background atmosphere or stopping depth, we simultaneously integrate the coupled momentum and deceleration equations, incorporating a generalized Coulomb logarithm valid across classical, quantum, and strongly magnetized regimes. This approach enables us to predict stopping lengths and compression profiles without imposing them by hand. Our results show that self-consistent atmospheric compression steepens the upper-layer density gradient and significantly modifies both the depth and shape of the braking region compared to fixed-density estimates. We find that stopping lengths range from ~0.01 cm in cool, high-Z atmospheres to ≥1 cm in hot, light compositions, with most kinetic energy dissipated in the bottom half of the braking layer where density and Coulomb coupling are largest. This coupling allows us to predict the stopping length and the braking-induced density gradient without imposing them by hand, yielding physically motivated vertical profiles that can be used directly in spectral and polarimetric calculations.

Theoretical model

Structure of Neutron Star Atmosphere.

The vertical structure of a neutron-star atmosphere in a subcritical X-ray pulsar is governed by the local gravitational field, the thermal pressure of the plasma, and the momentum exchange between the infalling accretion flow and the atmospheric layers. In this luminosity regime, radiation pressure is negligible, and the deceleration of the accretion flow is dominated by Coulomb collisions between the infalling protons and the atmospheric electrons. In several previous studies, the effect of plasma braking has been incorporated through a prescribed momentum-deposition profile, often parameterized as an exponential decay with optical depth, which provides a convenient phenomenological description of the stopping layer. While such prescriptions capture the localization of momentum transfer, they do not account for the feedback between flow deceleration and atmospheric compression, and therefore cannot predict the braking depth or density gradient self-consistently. In this work, we do not assume a predefined braking profile. Instead, we compute the deceleration of the accretion flow directly from Coulomb interactions, using a generalized Coulomb logarithm that remains valid in classical, quantum, and strongly magnetized regimes. The resulting braking rate depends explicitly on the local plasma density and temperature, which evolve as the atmosphere responds to momentum deposition. As a consequence, the velocity of the inflow and the atmospheric density structure are intrinsically coupled and must be solved simultaneously.

Under the assumption that the atmosphere adjusts quasi-statically to the applied force (justified by the short sound-crossing time compared to the accretion timescale), the atmospheric structure can be treated as being in local hydrostatic equilibrium, with an additional force term arising from momentum transfer by the decelerating flow. The vertical momentum balance can be written as

$$\frac{dP}{dz} = -\rho g + \frac{\dot{M}}{S}\frac{dv}{dz} \tag{1}$$

where the second term on the right-hand side represents the momentum transfer from the decelerating accretion flow to the atmospheric plasma. Here the gas pressure is given by $P = \rho kT/(\mu m_p)$, ρ is the local mass density, g is the gravitational acceleration, T is the local temperature, μ is the mean atomic weight, mp is the mass of a proton, S is the area of the hot spots at the NS surface, v is the local velocity of the accretion flow, and $\dot{M}$ is the mass accretion rate. The deceleration law for the flow is determined by Coulomb interactions between infalling protons and electrons in the atmosphere:

$$\frac{d}{dz}\left(\frac{m_p v^2}{z}\right) = -\frac{4\pi e^4 Z^2 n_e}{m_e v^2} ln\Lambda \tag{2}$$

where $\ln\Lambda$ is the Coulomb logarithm. While equation (2) is written for a non-magnetized plasma, it remains applicable in the magnetized case if $\ln\Lambda$ is replaced by the appropriate magnetized expression.[39] This approach allows the effects of strong magnetic fields to be incorporated through the microphysics of Coulomb scattering without modifying the macroscopic form of the braking equation. Introducing the dimensionless velocity $\beta = v/c$, the equation becomes:

$$\frac{d\beta}{dz} = 3.24 \times 10^{-4} \frac{\rho Z^2}{\beta^3} ln\Lambda \tag{3}$$

Combining equation (3) with the momentum balance equation (1), we obtain the density gradient in the atmosphere:

$$\frac{d\rho}{dz} = 0.1\frac{\mu g_{14}\rho}{T_{keV}} + 31.15\frac{\mu \dot{m}_6}{T_{keV}}\frac{d\beta}{dz} + \frac{\rho}{T_{keV}}\frac{dT_{keV}}{dz} \tag{4}$$

where $g_{14} \simeq 1.8 M_{1.4} R_6^{-2}$ is the gravitational acceleration in units of 10^{14} cm s^{-2}, $M_{1.4}$ is the NS mass in units of solar masses, R_6 is the NS radius in units of 10^6 cm, T_{keV} is the temperature in units of keV, and

$$\dot{m}_6 \simeq 25\, \dot{M}_{17}^{3/5} B_{12}^{1/2} \tag{5}$$

is the mass accretion rate per unit area measured in 10^6 g s^{-1} cm^{-2}, $\dot{M}_{17}$ is the mass accretion rate in units of 10^{17} g/s, and B_{12} is the surface magnetic field strength in units of 10^{12} G. In equation (4), the first term reflects hydrostatic compression by gravity, the second term represents additional atmospheric compression caused by momentum deposition from plasma braking, and the third term accounts for temperature gradients. In the absence of accretion ($\dot{m}_6 = 0$), equation (4) reduces to the standard hydrostatic equilibrium condition for a plane-parallel atmosphere. The density at the top of the atmosphere is set equal to the density of the accretion stream:

$$\rho_a = \frac{\dot{M}}{S_{vff}} \approx 2.5 \times 10^{-4} \frac{\dot{M}_{17} R_6^{1/2}}{S_{10} M_{1.4}^{1/2}} g cm^{-3} \tag{6}$$

which serves as the upper boundary condition for integrating equation (4). This condition assumes a smooth matching between the accretion stream and the atmospheric plasma; any finite transition layer would primarily affect the very top of the atmosphere and does not significantly alter the stopping depth or compression profile derived below.

The current model is designed for subcritical accretion onto strongly magnetized neutron stars. In this scenario, radiation pressure is minor compared to gravity, and Coulomb collisions primarily slow down the incoming plasma. The atmosphere is treated as plane-parallel and in steady state. This approach works well because the polar-cap footprint is small, and the sound-crossing time is short compared to the timescale of accretion. However, the model does not include radiative energy transport or temperature feedback in a self-consistent way; instead, it fixes the temperature profile to focus on the effects of plasma braking and atmospheric compression. Other effects are also ignored, in large part due to their relative insignificance. For example, multidimensional effects are neglected, as at the level of Coulomb collisions, the surrounding plasma remains almost perfectly uniform; likewise, magnetic field curvature is essentially negligible at such scales. General relativistic corrections are not taken into account, as these do not fundamentally change the electromagnetic interaction between particles. Given these assumptions, the model effectively shows the key relationship between Coulomb deceleration and atmospheric structure in the plasma-braking regime.

Coulomb Logarithm in Strong Magnetic Field and Low-Density Regime:

The Coulomb logarithm is generally defined as:

$$ln\Lambda = ln(\frac{b_{max}}{b_{min}}) \tag{7}$$

where b_{max} is the maximum impact parameter (typically the Debye length λ_D), and b_{min} is the minimum impact parameter, which requires more careful treatment in low-density, strongly magnetized, or quantum plasmas.

Thus, in this work, we use a generalized expression for the Coulomb logarithm, ln Λ, that remains valid in the presence of strong magnetic fields and extends to low-density plasma, where traditional approximations may break down. We adopt a synthesized formula that combines three physical scales: the classical distance of closest approach, the de Broglie wavelength, and the electron gyroradius (or Larmor radius). While no single prior publication has employed a similar generalization, the individual parameters are commonly seen. The classical distance of closest approach appears in non-relativistic Coulomb collision theory, while the de Broglie wavelength is used as the quantum lower bound on spatial resolution, and the Larmor radius limits perpendicular scattering in magnetized plasmas. We adopt the maximum of these three scales as bmin because at any given point, the collision process is regulated by whichever constraint is tightest, and taking the maximum ensures we do not underestimate the effective cutoff. The resulting formula provides a continuous and physically meaningful ln Λ across the full parameter space explored in this work.

$$ln\Lambda = ln(\frac{\lambda_D}{max\left\{\frac{e^2}{kT}, \frac{\hbar}{\sqrt{2m_e kT}}, \frac{\sqrt{2kT/m_e}}{\omega_{ce}}\right\}}) \tag{8}$$

where $\lambda_D = \sqrt{\frac{kT}{4\pi e^2 n_e}}$ is the Debye length, $\frac{e^2}{kT}$ is the classical distance of closest approach (i.e., the closest classical approach distance in the absence of quantum effects),[40] $\frac{\hbar}{\sqrt{2m_e kT}}$ is the de Broglie wavelength (this acts as the lower limit on spatial resolution due to wave-particle duality),[41] $\frac{\sqrt{2kT/m_e}}{\omega_{ce}}$ is the electron Larmor radius (in strongly magnetized plasma the motion perpendicular to the field becomes quantized, limiting the effective minimum distance for collisions),[42] $\omega_{ce}=eB/m_e c$ is the electron cyclotron frequency. Expression (8) ensures a physically meaningful lnΛ even when the plasma is near the collisionless regime or when quantum or magnetic effects dominate. For applications in neutron star envelopes and magnetized accretion flows, this form is consistent with previous works,[43-45] where magnetic field and quantum effects are included in transport coefficient calculations. This generalized expression is especially useful in simulations where plasma parameters evolve across regimes and a uniform treatment of lnΛ is desirable.

Throughout all models, the Coulomb logarithm lnΛ remains in the physically meaningful range $5 \lesssim \ln\Lambda \lesssim 20$, depending on the local density, temperature, and degree of magnetization. In the upper atmosphere, where the electron gyroradius exceeds both the classical and quantum scales, b_{min} is set by the Larmor radius, making the collisions effectively magnetized. Deeper in the atmosphere, as the density increases, the Debye length contracts and the plasma transitions to the classical regime in which e^2/kT dominates b_{min}. This evolution is naturally captured by the generalized expression of lnΛ and is reflected in the gradual variation of the braking efficiency with depth.

Numerical Model:

To model the deceleration of the accretion flow in the subcritical regime, we numerically integrate the coupled system of equations (3) and (4) governing the vertical evolution of the dimensionless velocity β(z) = v/c and the local mass density ρ(z). This approach simultaneously accounts for hydrostatic compression, plasma braking due to Coulomb collisions, and any temperature gradient effects. The Coulomb logarithm lnΛ appearing in equation (3) is evaluated self-consistently at each spatial step using the generalized expression given in equation (8). This form incorporates:

1) the classical distance of closest approach e^2/kT
2) the electron de Broglie wavelength $\frac{\hbar}{\sqrt{2m_e kT}}$, and
3) the electron gyroradius $v_\perp/\omega_{ce}$,

ensuring physical validity in classical, quantum, and strongly magnetized regimes. The maximum of these three scales is taken as b_{min}, while the Debye length λ_D is used as bmax, giving $\ln\Lambda = \ln(\lambda_D/b_{min})$. The numerical domain begins at the top of the atmosphere, where the inflow velocity is set to β_i (typically corresponding to the local free-fall speed), and the density is fixed to the accretion stream value ρ_a from equation (6). Instead of prescribing a single value of μ, we specify the composition of the atmospheric plasma through the ion mass number A

DOI: 10.36838/IJHSR820.7

and charge Z. The corresponding mean molecular weights are computed as $\mu = \frac{A}{Z+1}$, $\mu_e = \frac{A}{Z}$, representing the mean mass per particle (ions + electrons) and per electron, respectively. These quantities enter the pressure term and the Coulomb logarithm in equations (3) and (4). While the temperature can be held constant for simplicity, the formulation allows for an arbitrary T(z) profile, in which case the third term in equation (4) is non-zero. We employ an explicit finite-difference scheme in the physical height x with a fixed step Δx small enough to resolve the steepest expected gradients. At each step:

1) Velocity update — We compute $\frac{d\beta}{dz}$ from equation (3) using the current ρ and lnΛ, and update β via an Euler step.

2) Density update — With the new β, we then calculate $\frac{d\rho}{dz}$ from equation (4), which includes contributions from gravitational compression, momentum exchange from braking, and any imposed temperature gradient. The new density is then updated to the next step.

3) Termination — The integration stops when $\beta \to 0$ or becomes negative, which we define as the total stopping length, x_{stop}.

The output includes the full spatial profiles β(x) and ρ(x), as well as diagnostic quantities such as the local braking rate, the evolution of lnΛ, and the fraction of kinetic energy dissipated at each depth. To highlight the influence of composition, we present results for different choices of (A, Z), including hydrogen and carbon atmospheres, which correspond to μ = 1/2 and μ = 12/7, respectively. For clarity, we also show constant-ρ runs (to isolate kinematic effects) alongside full coupled β–ρ integrations (to capture the interplay between flow deceleration and atmospheric compression).

Numerical Results

In this section, we present the results of the coupled integration of the velocity and density structure of the neutron-star atmosphere for a range of temperatures and chemical compositions. Unless stated otherwise, all models assume a surface magnetic field $B = 10^{12}$ G, gravitational acceleration $g_{14} = 2.2$, initial velocity $\beta_i = 0.5$, and initial mass density $\rho_0 = 10$ g cm^{-3}. The numerical step was fixed at $\Delta x = 10^{-3}$ cm, sufficient to resolve the shortest braking lengths obtained in any model.

Temperature Influencing Plasma Brake:

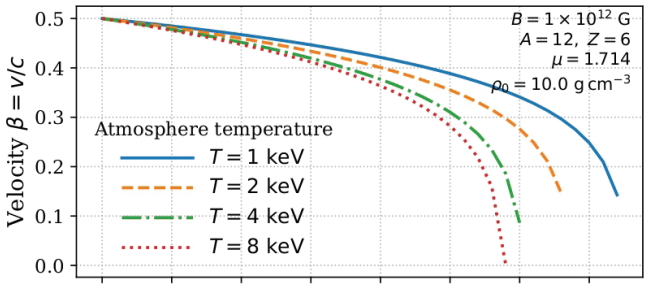


Figure 1: Velocity β(x) profiles for a temperature scan of a fully ionized carbon atmosphere (A = 12, Z = 6, μ = 12/7) with evolving density. Assume $B = 10^{12}$ G, surface gravity $g = 2.2 \times 10^{14}$ cm s^{-2}, initial density $\rho_0 = 10$ g cm^{-3}, and initial velocity $\beta_i = 0.5$.

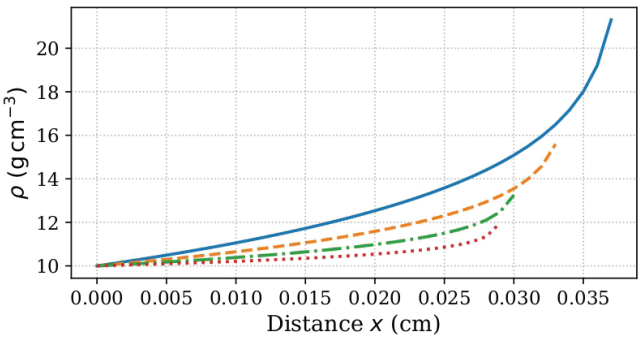


Figure 2: Density ρ(x) profiles with the same initial parameters. The curves show solutions of the plasma-braking equations for constant temperatures T = 1, 2, 4, and 8 keV. A stronger thermal pressure support at higher T reduces the deceleration rate and lowers the density gradient, producing a more extended braking region.

Figures 1 and 2 show the velocity β(x) and density ρ(x) profiles for a set of constant-temperature atmospheres with T = 1–8 keV. These models use a fully ionized carbon composition with (A, Z) = (12, 6) and μ = 12/7. The stopping length is found to depend sensitively on the thermal state of the atmosphere. At low temperature (T = 1–2 keV), the Coulomb drag is low due to the small electron thermal velocity and hence small Coulomb Logarithm, and the inflowing protons decelerate slowly, producing a broad braking zone of thickness Δx ~ a few millimeters. As the temperature increases, the electron thermal speed grows, and the Coulomb logarithm increases moderately, leading to a more effective drag force. This behavior results in a progressively more compact braking region and a steeper velocity gradient. The atmospheric response follows the same trend. At higher temperatures, the compression is steeper, and the peak density attained before complete stopping is correspondingly larger. These temperature-dependent density gradients are essential for determining the location of the vacuum resonance and the degree of adiabatic mode conversion, as discussed in the Conclusion.

Physically, increasing the temperature increases the Debye length, which in turn increases the Coulomb Logarithm. This causes the Coulomb drag to be more effective, resulting in the flow losing momentum faster and hence stopping over a shorter distance.

Chemical Composition Influencing Plasma Brake:

To quantify the dependence on ion mass and charge, we compare in Figure 3 hydrogen and carbon atmospheres at a fixed temperature T = 2 keV. These compositions correspond to (A, Z) = (1, 1) with μ = 1/2 and (A, Z) = (12, 6) with μ = 12/7, respectively. The braking behavior differs strongly between the two cases. In a hydrogen atmosphere, the Coulomb drag is intrinsically weaker because the collision term scales as Z^2 and because the larger number of particles per unit mass (smaller μ) reduces the atmospheric inertia. Consequently, the protons penetrate deeper before being stopped, and the resulting density enhancement is relatively modest. The density gradient is shallow, and the velocity decreases gradually over a distance of order ~ 1 cm. In contrast, a carbon atmosphere provides both a higher collision rate (Z2 = 36) and a larger mean par-

ticle mass (μ = 12/7). The combination of stronger Coulomb coupling and a more inert medium causes rapid deceleration within ~ 0.1 cm. The atmospheric density correspondingly increases over a smaller spatial scale, producing a much steeper density gradient in the braking layer. This behavior reflects the combined effects of stronger Coulomb coupling and increased atmospheric inertia in high-Z compositions: a larger charge enhances the collision rate, while a higher mean molecular weight reduces the ability of the atmosphere to redistribute the deposited momentum, leading to rapid deceleration and strong local compression.

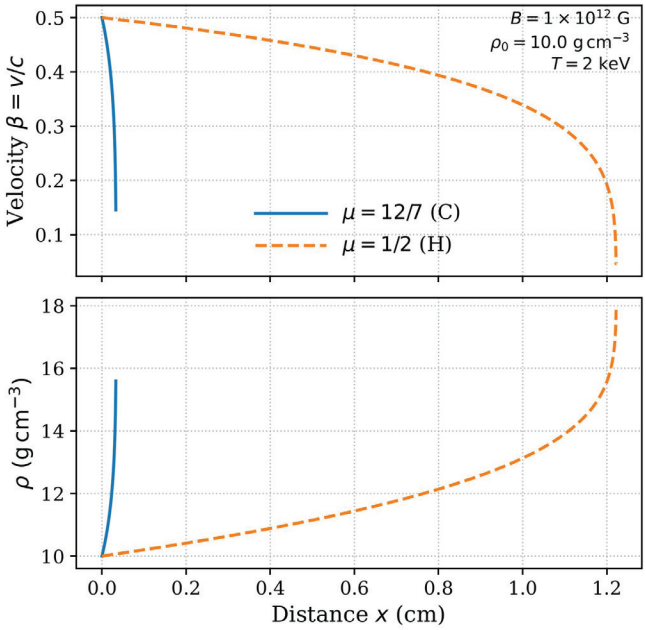


Figure 3: Comparison of velocity β(x) and density ρ(x) profiles for two fully ionized compositions—hydrogen (A = 1, Z = 1, μ = 1/2) and carbon (A = 12, Z = 6, μ = 12/7)—at a fixed temperature T = 2 keV. The differing mean molecular weights and charge states modify both the Coulomb drag (through Z^2) and the hydrostatic compression term (through μ), leading to distinct braking lengths and density structures. All other parameters are the same as in Figures 1 and 2.

Stopping Lengths and Energy Dissipation:

The total stopping length xstop ranges from ~0.01 cm in cool, high-Z atmospheres to ≳ 1 cm in light, hot compositions. In all cases, most of the kinetic energy is dissipated in the bottom half of the braking layer, where both the density and the Coulomb coupling are largest. The fraction of the incident kinetic energy converted into local heating within each depth interval can be computed directly from the drag force and is naturally stored as part of the numerical diagnostics, although we defer a detailed analysis of this energy deposition profile to future work. Overall, the numerical results demonstrate that both temperature and composition play decisive roles in shaping the vertical structure of the atmosphere in sub-critical XRPs. The steepening or smoothing of the density gradient has a direct influence on radiative transfer, vacuum resonance crossing, and ultimately the polarization degree of the emergent X-ray radiation.

Atmospheric Compression by Coulomb Braking:

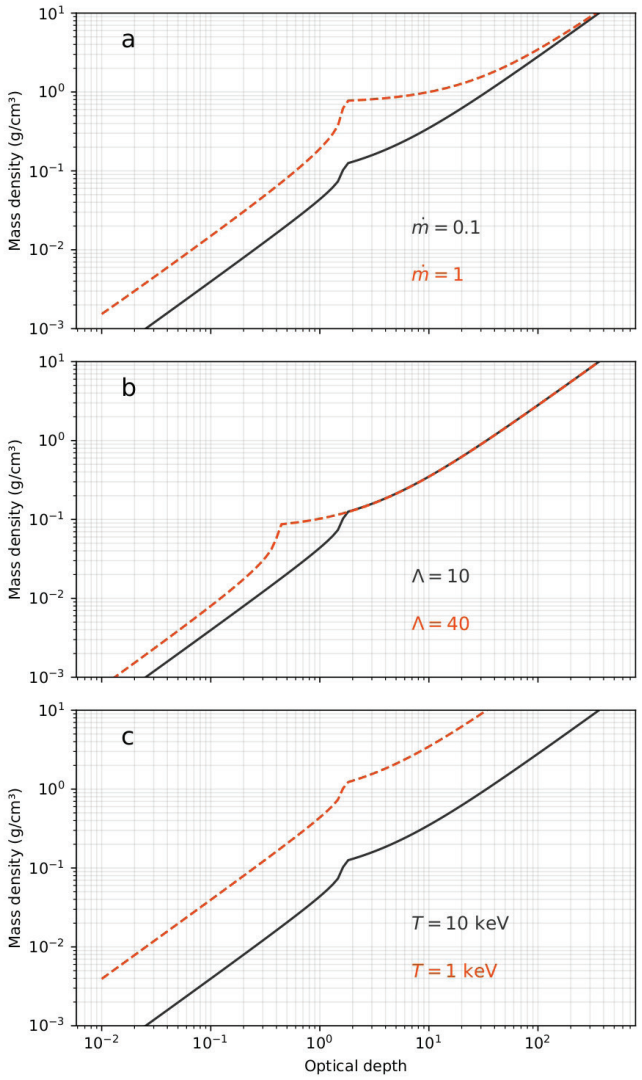


Figure 4: Mass density profiles ρ(τ) as a function of optical depth for a set of atmospheric models illustrating the influence of (a) the mass accretion rate per unit area, (b) the Coulomb logarithm, and (c) the atmospheric temperature. In all panels, the horizontal axis shows the optical depth with respect to non-magnetic Thomson scattering, while the vertical axis shows the local mass density.

Figure 4 illustrates how Coulomb braking of the accretion flow imprints a characteristic vertical structure on the atmosphere of a strongly magnetized NS. The horizontal axis in all panels shows the optical depth with respect to non-magnetic Thomson scattering, while the vertical axis gives the local mass density. In each model, the inflowing plasma deposits its momentum within a finite stopping layer, causing a distinct compression feature: the density rises steeply once the protons begin to decelerate efficiently, producing a pronounced "knee" in the ρ(τ) profile. The location and sharpness of this feature directly reflect the depth at which the flow comes to rest. Panel (a) shows the dependence of the atmospheric response on the mass accretion rate per unit area. The black curve corresponds to a ten times lower inflow rate than the red curve, and the resulting compression is accordingly weaker. A larger accretion flux transfers more momentum to the atmosphere

DOI: 10.36838/IJHSR820.7

and forces the braking region to greater density and optical depth, yielding a systematically higher density throughout the upper layers. This behavior emphasizes that even sub-critical accretion can substantially reshape the stratification of the atmosphere when the inflow is sufficiently intense. Panel (b) shows how the result changes with the value of the Coulomb logarithm, which quantifies the efficiency of collisional momentum transfer. Increasing $\ln\Lambda$—from 10 in the black curve to 40 in the red curve—strengthens the deceleration of the flow. As a consequence, the inflowing protons are stopped at smaller optical depths, producing a sharper and more elevated density enhancement in the upper atmosphere. At larger depths, however, both models converge to the same hydrostatic gradient, indicating that the braking mainly alters the structure above the stopping layer while leaving the deep atmosphere essentially unaffected. Panel (c) explores the influence of the atmospheric temperature. A cooler atmosphere (red curve, T = 1 keV) has a higher density at a given optical depth due to reduced thermal pressure support, and therefore experiences more efficient Coulomb braking. Although both the hot and cool models exhibit the same qualitative signature of a density jump associated with the termination of the inflow, the entire density profile is shifted upward in the colder case. This systematic offset reflects the competing roles of thermal support and momentum deposition in determining the vertical structure. Overall, Figure 4 highlights that the interplay between accretion dynamics and atmospheric microphysics naturally generates strong compression of the upper layers of the neutron-star atmosphere. The depth and sharpness of the braking-induced density gradient depend sensitively on the accretion rate, the Coulomb coupling strength, and the thermal state of the plasma. These variations modify the local optical depth profile and are therefore expected to influence radiative transfer, vacuum-resonance behavior, and ultimately the observed polarization properties of sub-critical X-ray pulsars.

■ Conclusion

In this work, we have investigated the deceleration of accretion flows in sub-critical X-ray pulsars, focusing on the regime in which radiation pressure is negligible and Coulomb collisions in a strongly magnetized neutron star atmosphere dominate the braking process. Using a self-consistent numerical integration of the coupled velocity and density equations, we have quantified how the accretion flow – atmosphere interaction shapes the vertical structure of the accretion flow base. Our results demonstrate that plasma braking naturally produces a strongly compressed upper atmosphere, with a characteristic steepening of the density gradient once the infalling protons begin to lose momentum efficiently. This compression feature provides a clear dynamical signature of the transition between the decelerating flow and the underlying hydrostatic atmosphere. The location, amplitude, and sharpness of this compression feature depend sensitively on the local plasma parameters. Higher mass accretion rates lead to more substantial atmospheric compression and shift the braking region to larger optical depths (see Figure 4a). Variations in the Coulomb logarithm alter the efficiency of collisional momentum transfer, modifying the stopping depth and the structure of the upper layers (see Figure 4b). The thermal state of the atmosphere likewise plays a key role (see Figure 4c): cooler plasmas brake the inflow more rapidly and yield higher peak densities, while hotter atmospheres produce more extended and gradual deceleration zones. These findings have important implications for radiative transfer of polarized X-rays in sub-critical XRPs. The density gradients produced by Coulomb braking regulate the local optical depth and determine whether the normal modes of radiation propagate adiabatically or undergo partial mode conversion near the vacuum resonance.[2] In particular, because the adiabaticity of mode propagation depends sensitively on the density scale height, the braking-induced modification of $\rho(z)$ can change the probability of non-adiabatic switching between the O- and X-modes near the vacuum resonance, thereby affecting the polarization state of the emerging radiation. The braking-induced stratification, therefore, represents a key physical ingredient in modeling the emergent spectra and polarization signatures of accreting neutron stars. The framework developed here provides a physically consistent description of the braking layer and its dependence on accretion rate, composition, and temperature. The resulting density profiles can be used directly as input for spectral, polarimetric, and atmosphere-modeling calculations. We emphasize that the main contribution of this work is to provide self-consistent vertical structure profiles in the plasma braking regime, which can be straightforwardly coupled to existing radiative-transfer and polarization pipelines. Future work will extend the present approach by incorporating nonconstant temperature profiles, radiative feedback, and multidimensional effects, enabling a more complete description of the accretion footprint in both sub-critical and transitional luminosity regimes.

Future Work:

Several extensions of this work could provide a more complete picture of sub-critical accretion.

First, incorporating self-consistent temperature profiles, calculated from energy deposition and radiative cooling, would replace the isothermal approximation. This would capture the thermal feedback between deceleration and atmospheric structure, which was assumed to be constant in this research.

Second, including radiative feedback, especially through cyclotron cooling and free-free emission, would allow us to model the shift from sub-critical to near-critical states. In these states, radiation pressure starts to affect the dynamics.

Third, extending the model to multiple spatial dimensions, either 2D or 3D, would help us study the lateral spread of the accretion flow. It would also allow us to examine the formation of uneven density structures within the hot spot and how magnetic field geometry influences vertical layering.

■ Acknowledgments

I want to express my deepest gratitude to my mentor, Dr. Alexander Mushtukov, for their invaluable guidance, encouragement, and support throughout this research project. Their

insightful feedback and expertise were instrumental in shaping the methodology and the final research paper.

References

1. Mushtukov, A.; Tsygankov, S. Accreting strongly magnetized neutron stars: X-ray Pulsars. arXiv.org. https://arxiv.org/abs/2204.14185 (accessed 2025-12-21).
2. Harding, A. K.; Lai, D. Physics of Strongly Magnetized Neutron Stars. Reports on Progress in Physics 2006, 69 (9), 2631–2708. https://doi.org/10.1088/0034-4885/69/9/r03.
3. Illarionov, A. F.; Sunyaev, R. A. Why the Number of Galactic X-ray Stars Is so Small? Astronomy and Astrophysics, Vol. 39, p. 185. https://adsabs.harvard.edu/full/1975A%26A....39..185I
4. Romanova, M. M.; Ustyugova, G. V.; Koldoba, A. V.; Lovelace, E. Propeller-Driven Outflows and Disk Oscillations. The Astrophysical Journal 2005, 635 (2), L165–L168. https://doi.org/10.1086/499560.
5. Tsygankov, S. S.; Lutovinov, A. A.; Doroshenko, V.; Mushtukov, A. A.; V. Suleimanov; J. Poutanen. Propeller Effect in Two Brightest Transient X-Ray Pulsars: 4U 0115+63 and v 0332+53. Astronomy and Astrophysics 2016, 593, A16–A16. https://doi.org/10.1051/0004-6361/201628236.
6. Tsygankov, S. S.; Mushtukov, A. A.; Suleimanov, V. F.; Doroshenko, V.; Abolmasov, P. K.; Lutovinov, A. A.; J. Poutanen. Stable Accretion from a Cold Disc in Highly Magnetized Neutron Stars. Astronomy and Astrophysics 2017, 608, A17–A17. https://doi.org/10.1051/0004-6361/201630248.
7. Fabrika, S. N.; Atapin, K. E.; Vinokurov, A. S.; Sholukhova, O. N. Ultraluminous X-Ray Sources. Astrophysical Bulletin 2021, 76 (1), 6–38. https://doi.org/10.1134/s1990341321010077.
8. King, A.; Lasota, J.-P.; Middleton, M. Ultraluminous X-Ray Sources. New Astronomy Reviews 2023, 96, 101672. https://doi.org/10.1016/j.newar.2022.101672.
9. Zel'dovich, Y. B.; Shakura, N. I. X-Ray Emission Accompanying the Accretion of Gas by a Neutron Star. Soviet Astronomy. Vol. 13, p. 175. https://adsabs.harvard.edu/full/1969SvA....13..175Z
10. Nelson, R. W.; John; Salpeter, E. E.; Wasserman, I. A Potential Cyclotron Line Signature in Low-Luminosity X-Ray Sources. The Astrophysical Journal 1995, 438, L99–L99. https://doi.org/10.1086/187725.
11. Basko, M. M.; Sunyaev, R. A. The Limiting Luminosity of Accreting Neutron Stars with Magnetic Fields. Monthly Notices of the Royal Astronomical Society 1976, 175 (2), 395–417. https://doi.org/10.1093/mnras/175.2.395.
12. Mushtukov, A. A.; Suleimanov, V. F.; Tsygankov, S. S.; Juri Poutanen. The Critical Accretion Luminosity for Magnetized Neutron Stars. 2015, 447 (2), 1847–1856. https://doi.org/10.1093/mnras/stu2484.
13. Sheng, X.; Zhang, L.; Blaes, O.; Jiang, Y.-F. Dynamical Effects of Magnetic Opacity in Neutron Star Accretion Columns. Monthly Notices of the Royal Astronomical Society 2023, 524 (2), 2431–2445. https://doi.org/10.1093/mnras/stad2043.
14. Tsygankov, S. S.; Rouco Escorial, A.; Suleimanov, V. F.; Mushtukov, A. A.; Doroshenko, V.; Lutovinov, A. A.; Wijnands, R.; Juri Poutanen. Dramatic Spectral Transition of X-Ray Pulsar GX 304−1 in Low Luminous State. Monthly Notices of the Royal Astronomical Society Letters 2018, 483 (1), L144–L148. https://doi.org/10.1093/mnrasl/sly236.
15. Tsygankov, S. S.; Doroshenko, V.; Mushtukov, A. A.; V. Suleimanov; A. Lutovinov; Juri Poutanen. Cyclotron Emission, Absorption, and the Two Faces of X-Ray Pulsar a 0535+262. Monthly Notices of the Royal Astronomical Society: Letters 2019, 487 (1), L30–L34. https://doi.org/10.1093/mnrasl/slz079.
16. A. Lutovinov; S. Tsygankov; S. Molkov; Doroshenko, V.; A. Mushtukov; V. Arefiev; I. Lapshov; Tkachenko, A.; M. Pavlinsky. SRG/ART-XC and NuSTAR Observations of the X-Ray Pulsar GRO J1008–57 in the Lowest Luminosity State. The Astrophysical Journal 2021, 912 (1), 17–17. https://doi.org/10.3847/1538-4357/abec43.
17. Mushtukov, A. A.; V. Suleimanov; S. Tsygankov, S. P. Zwart. Spectrum formation in X-ray pulsars at very low mass accretion rate: Monte Carlo approach, Monthly Notices of the Royal Astronomical Society, Volume 503, Issue 4, June 2021, Pages 5193–5203, https://doi.org/10.1093/mnras/stab811
18. Sokolova-Lapa, E.; Gornostaev, M.; Wilms, J.; Ballhausen, R.; Falkner, S.; Postnov, K.; Thalhammer, P. X-Ray Emission from Magnetized Neutron Star Atmospheres at Low Mass-Accretion Rates. I. Phase-Averaged Spectrum. Astronomy & Astrophysics 2021. https://doi.org/10.1051/0004-6361/202040228.
19. Becker, P. A.; Wolff, M. T. Thermal and Bulk Comptonization in Accretion-Powered X-Ray Pulsars. The Astrophysical Journal 2007, 654 (1), 435–457. https://doi.org/10.1086/509108.
20. Daugherty, J. K.; Harding, A. K. Compton Scattering in Strong Magnetic Fields. The Astrophysical Journal 1986, 309, 362. https://doi.org/10.1086/164608.
21. Lai, D.; Wynn. Transfer of Polarized Radiation in Strongly Magnetized Plasmas and Thermal Emission from Magnetars: Effect of Vacuum Polarization. Astrophysical journal/The Astrophysical journal 2003, 588 (2), 962–974. https://doi.org/10.1086/374334.
22. Zheleznyakov, V. V. Electromagnetic waves in cosmic plasma. Generation and propagation. Electromagnetic waves in cosmic plasma. Generation and propagation. https://ui.adsabs.harvard.edu/abs/1977ewcp.book.....Z/abstract (accessed 2025-12-21).
23. Meszaros, P.; Nagel, W. X-Ray Pulsar Models. I - Angle-Dependent Cyclotron Line Formation and Comptonization. The Astrophysical Journal 1985, 298, 147–147. https://doi.org/10.1086/163594.
24. Meszaros, P.; Nagel, W. X-Ray Pulsar Models. II - Comptonized Spectra and Pulse Shapes. The Astrophysical Journal 1985, 299, 138–138. https://doi.org/10.1086/163687.
25. Caiazzo, I.; Heyl, J. Polarization of Accreting X-Ray Pulsars. I. A New Model. Monthly Notices of the Royal Astronomical Society 2020, 501 (1), 109–128. https://doi.org/10.1093/mnras/staa3428.
26. M. C. Weisskopf; P. Soffitta; L. Baldini; B. D. Ramsey; S. L. O'Dell; R. W. Romani; G. Matt; W. D. Deininger; W. H. Baumgartner; R. Bellazzini; E. Costa; J. J. Kolodziejczak; L. Latronico; H. L. Marshall; F. Muleri; S. D. Bongiorno; A. Tennant; N. Bucciantini; M. Dovciak; F. Marin; A. Marscher; J. Poutanen; P. Slane; R. Turolla; W. Kalinowski; A. Di Marco; S. Fabiani; M. Minuti; F. La Monaca; M. Pinchera; J. Rankin; C. Sgrò; A. Trois; F. Xie; C. Alexander; D. Z. Allen; F. Amici; J. Andersen; A. Antonelli; S. Antoniak; P. Attiná; M. Barbanera; M. Bachetti; R. M. Baggett; J. Bladt; A. Brez; R. Bonino; C. Boree; F. Borotto; S. Breeding; D. Brienza; H. K. Bygott; C. Caporale; C. Cardelli; R. Carpentiero; S. Castellano; M. Castronuovo; L. Cavalli; E. Cavazzuti; M. Ceccanti; M. Centrone; S. Citraro; F. D'Amico; E. D'Alba; L. Di Gesu; E. Del Monte; K. L. Dietz; N. Di Lalla; G. Di Persio; D. Dolan; I. Donnarumma; Y. Evangelista; K. Ferrant; R. Ferrazzoli; M. Ferrie; J. Footdale; B. Forsyth; M. Foster; B. Garelick; S. Gunji; E. Gurnee; M. Head; G. Hibbard; S. Johnson; E. Kelly; K. Kilaru; C. Lefevre; S. Le Roy; P. Loffredo; P. Lorenzi; L. Lucchesi; T. Maddox; G. Magazzu; S. Maldera; A. Manfreda; E. Mangraviti; M. Marengo; A. Marrocchesi; F. Massaro; D. Mauger; J. McCracken; M. McEachen; R. Mize; P. Mereu; S. Mitchell; I. Mitsuishi; A. Morbidini; F. Mosti; H. Nasimi; B. Negri; M. Negro; T. Nguyen; I. Nitschke; A. Nuti;

DOI: 10.36838/IJHSR820.7

M. Onizuka; C. Oppedisano; L. Orsini; D. Osborne; R. Pacheco; A. Paggi; W. Painter; S. D. Pavelitz; C. Pentz; R. Piazzolla; M. Perri; M. Pesce-Rollins; C. Peterson; M. Pilia; A. Profeti; S. Puccetti; J. Ranganathan; A. Ratheesh; L. Reedy; N. Root; A. Rubini; S. Ruswick; J. Sanchez; P. Sarra; F. Santoli; E. Scalise; A. Sciortino; C. Schroeder; T. Seek; K. Sosdian; G. Spandre; C. O. Speegle; T. Tamagawa; M. Tardiola; A. Tobia; N. E. Thomas; R. Valerie; M. Vimercati; A. L. Walden; B. Weddendorf; J. Wedmore; D. Welch; D. Zanetti; F. Zanetti. Imaging X-ray Polarimetry Explorer: prelaunch Journal of Astronomical Telescopes, Instruments, and Systems, Vol. 8, Issue 2, 026002 (April 2022). https://doi.org/10.1117/1.JATIS.8.2.026002

27. Doroshenko, V.; Juri Poutanen; Tsygankov, S. S.; Suleimanov, V. F.; Bachetti, M.; Caiazzo, I.; Costa, E.; Marco, A. D.; Heyl, J.; Fabio La Monaca; Fabio Muleri; Mushtukov, A. A.; Pavlov, G. G.; Ramsey, B. D.; Rankin, J.; Santangelo, A.; Paolo Soffitta; Rüdiger Staubert; Weisskopf, M. C.; Zane, S. Determination of X-Ray Pulsar Geometry with IXPE Polarimetry. Nature Astronomy 2022, 6 (12), 1433–1443. https://doi.org/10.1038/s41550-022-01799-5.
28. Tsygankov, S. S.; Doroshenko, V.; Juri Poutanen; Heyl, J.; Mushtukov, A. A.; Caiazzo, I.; Marco, A. D.; Forsblom, S. V.; González-Caniulef, D.; Moritz Klawin; Fabio La Monaca; Malacaria, C.; Marshall, H. L.; Fabio Muleri; Ng, M.; Suleimanov, V. F.; Sunyaev, R. A.; Turolla, R.; Agudo, I.; Antonelli, L. A. The X-Ray Polarimetry View of the Accreting Pulsar Cen X-3. The Astrophysical Journal Letters 2022, 941 (1), L14–L14. https://doi.org/10.3847/2041-8213/aca486.
29. Tsygankov, S. S.; Doroshenko, V. T.; Mushtukov, A. A.; Juri Poutanen; Alessandro Di Marco. X-Ray Pulsar GRO J1008–57 as an Orthogonal Rotator. Astronomy and Astrophysics 2023, 675, A48–A48. https://doi.org/10.1051/0004-6361/202346134.
30. Mushtukov, A. A.; Tsygankov, S. S.; J Poutanen; Doroshenko, V.; A Salganik; Costa, E.; Marco, A. D.; Heyl, J.; F La Monaca; Lutovinov, A. A.; Mereminsky, I. A.; A Papitto; Semena, A. N.; Shtykovsky, A. E.; Suleimanov, V. F.; Forsblom, S. V.; D González-Caniulef; C Malacaria; Sunyaev, R. A.; Agudo, I. X-Ray Polarimetry of X-Ray Pulsar X Persei: Another Orthogonal Rotator? Monthly Notices of the Royal Astronomical Society 2023, 524 (2), 2004–2014. https://doi.org/10.1093/mnras/stad1961.
31. Forsblom, S. V.; Juri Poutanen; Tsygankov, S. S.; Bachetti, M.; Marco, A. D.; Doroshenko, V.; Heyl, J.; Fabio La Monaca; Malacaria, C.; Marshall, H. L.; Fabio Muleri; Mushtukov, A. A.; Pilia, M.; Rogantini, D.; Suleimanov, V. F.; Taverna, R.; Xie, F.; Agudo, I.; Antonelli, L. A.; Baldini, L. IXPE Observations of the Quintessential Wind-Accreting X-Ray Pulsar Vela X-1. The Astrophysical Journal Letters 2023, 947 (2), L20–L20. https://doi.org/10.3847/2041-8213/acc391.
32. Suleimanov, V. F.; Forsblom, S. V.; S.S. X-Ray Polarimetry of the Accreting Pulsar GX 301–2. Astronomy & Astrophysics 2023, 678, A119–A119. https://doi.org/10.1051/0004-6361/202346994.
33. Poutanen, J.; Tsygankov, S. S.; Forsblom, S. V. X-Ray Polarimetry of X-Ray Pulsars. Galaxies 2024, 12 (4), 46. https://doi.org/10.3390/galaxies12040046.
34. Lai, D.; Wynn. Resonant Conversion of Photon Modes due to Vacuum Polarization in a Magnetized Plasma: Implications for X-Ray Emission from Magnetars. Astrophysical journal/The Astrophysical journal 2002, 566 (1), 373–377. https://doi.org/10.1086/338074.
35. Adler, S. L. Photon Splitting and Photon Dispersion in a Strong Magnetic Field. Annals of Physics 1971, 67 (2), 599–647. https://doi.org/10.1016/0003-4916(71)90154-0.
36. Tsai, W.; Erber, T. Propagation of Photons in Homogeneous Magnetic Fields: Index of Refraction. Physical Review. D. Particles, fields, gravitation, and cosmology/Physical Review. D. Particles and fields 1975, 12 (4), 1132–1137. https://doi.org/10.1103/physrevd.12.1132.
37. Mészáros, P.; Ventura, J. Vacuum Polarization Effects on Radiative Opacities in a Strong Magnetic Field. Physical Review D 1979, 19 (12), 3565–3575. https://doi.org/10.1103/physrevd.19.3565.
38. Pavlov,. Vacuum Polarization by a Magnetic Field and Its Astrophysical Manifestations. Astrophysics and Space Physics Reviews 2025, 3, 197.
39. Nelson, R. W.; Salpeter, E. E.; Wasserman, I. Nonthermal Cyclotron Emission from Low-Luminosity Accretion onto Magnetic Neutron Stars. The Astrophysical Journal 1993, 418, 874. https://doi.org/10.1086/173445.
40. Spitzer, L. Physics of Fully Ionized Gases. Physics of Fully Ionized Gases. https://ui.adsabs.harvard.edu/abs/1962pfig.book.....S/abstract (accessed 2025-12-22).
41. Ichimaru, S. Basic principles of plasma physics, a statistical approach. Basic principles of plasma physics. https://ui.adsabs.harvard.edu/abs/1973bppp.book.....I/abstract (accessed 2025-12-22).
42. Yakovlev D. G.; Urpin V. A. Thermal and Electrical Conductivity in White Dwarfs and Neutron Stars. Soviet Astronomy 1980, 24, 303.
43. A. Y. Potekhin; D. A. Baiko; P. Haensel; D. G. Yakovlev Transport properties of degenerate electrons in neutron star envelopes and white dwarf cores. arXiv.org. https://arxiv.org/abs/astro-ph/9903127 (accessed 2025-12-22).
44. Potekhin, A. Y.; Chabrier, G. Equation of State for Magnetized Coulomb Plasmas. Astronomy and Astrophysics 2013, 550, A43–A43. https://doi.org/10.1051/0004-6361/201220082.
45. Potekhin, A. Y.; Chabrier, G. Magnetic Neutron Star Cooling and Microphysics. Astronomy and Astrophysics 2018, 609, A74–A74. https://doi.org/10.1051/0004-6361/201731866.

Author

Chen Xingyu is a 12th grader studying at Raffles Institution, Singapore. He combines a love for mathematics and physics with a collaborative mindset. Driven by a passion for bridging the gap between abstract theoretical models and observed data, he very much wants to pursue a future in physics and engineering.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

The Effect of Environmental Regulation Stringency on Productivity and Innovation: Testing the Porter Hypothesis

Chris Z. Kumala

Rangitoto College, 564 East Coast Road, Mairangi Bay, Auckland, 0630, New Zealand; Chriskumala88@outlook.com
Mentor: Dr. Ceyhun Elgin

ABSTRACT: Decisions made by the manufacturing industry are often presented as a tradeoff between efficiency and environmental responsibility. Contrary to this, Porter's Hypothesis (PH) argues that well-designed environmental regulation can stimulate innovation, possibly increasing productivity. We explore the relationship between the stringency of environmental regulation and short-term changes in productivity within the manufacturing sector. This paper hypothesizes that stringent environmental regulation leads to no increase in productivity in regulated manufacturing sectors; rather, inventive activity may increase, in line with the Weak PH. Using OECD productivity and EU patent data, combined with EU Emissions Trading System (ETS) emissions and allowance data across OECD countries, this study uses panel data at the country–industry level from 1998 to 2010. Baseline regressions employ a two-way fixed effects model with country–industry linear trends to estimate the association between lagged log regulatory stringency and productivity and innovation, alongside event studies around the ETS introduction. The results show that lagged log stringency is not statistically associated with short-run changes in productivity or innovation, with coefficients close to zero across periods. Overall, the findings show limited support for both the Weak and Strong PH in the short run and suggest policymakers should not expect instant gains following environmental regulation.

KEYWORDS: Economics, Environmental Economics, Climate Policy, Porter's Hypothesis, EU Emissions Trading System.

Introduction

The neoclassical view holds that stringent environmental regulation impedes innovation and thereby growth on the sector level. Proponents of stringent environmental regulation counter this view, arguing that increasing the stringency of environmental regulation provides just cause for firms to innovate in order to reduce new compliance costs. Under this assumption, sector-wide productivity increases while environmental spillover costs decrease. This process is based on Schumpeter's idea of Creative Destruction[1], whereby new processes make old processes obsolete, resulting in more productive technologies. Consistent with endogenous growth frameworks such as Romer[2], regulatory pressure may therefore affect short-run productivity through induced technological change. Porter and van der Linde formalized this mechanism[3] with their theory, coining Porter's Hypothesis. Specifically, they argued that markets with asymmetric information, organizational inertia, and control problems create market imperfections. When environmental regulation is introduced, firms are pushed to react by exploring overlooked paths for technological improvement.

Further literature has created two distinct ways the PH can act at the firm level. Jaffe *et al.* are commonly credited[4] with articulating the "Weak" and "Strong" PH. First, the "Weak" PH claims that well-designed environmental regulation induces innovation accompanied by a limited or unclear impact on sector-level productivity. Second, the "Strong" PH postulates that well-designed regulation leads to innovation in technologies that improve a firm's productivity, thereby causing a net positive effect on sector-level productivity. While the "Weak" PH is widely accepted as true within the medium and long-term, recent scholarship has suggested that the effects of environmental regulation differ substantially across sectors, policy instruments, and time periods. Dechezleprêtre and Sato[5] supported the notion that environmental regulation can stimulate green innovation. Yet, increases in competitiveness and productivity are subject to policy design and adjustment periods. Similarly, Ambec *et al.*[6] suggested that regulation must be sufficiently stringent for the Porter-style effects to occur. Therefore, the regulation design is a likely factor in determining whether innovation offsets compliance costs in the short run. Whether the "Strong" PH is backed by empirical data is more heavily contested.

This paper assesses whether stringent environmental regulation leads to short-term productivity gains in the manufacturing sector. This question has particular importance even beyond environmental economics, as governments increasingly adopt environmental regulation; understanding the effects of such policy is important for policymakers, so informed decisions can be made. Addressing it requires engaging with empirical literature that tests Porter's Hypothesis and evaluates whether innovation caused by stringent environmental regulation translates into concrete productivity gains. Nonetheless, Porter's initial hypothesis was rather constrained as it was primarily anecdotal and based on limited empirical evidence. Hence, contemporary scholarship has sought to use data to reveal the true relationship between the stringency of environmental regulation and productivity gains. These studies have been largely focused on the United States and other OECD economies, with emphasis on the medium to long-term effects on energy-intensive manufacturing sectors.

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.8

However, recent papers have reached heterogeneous conclusions. Earlier studies, particularly those focusing on the United States, found that compliance costs associated with pollution abatement imposed net costs on firms, with limited evidence of productivity-enhancing innovation. Gray and Shadbegian[7] examined firm-level data from US manufacturing industries, concluding that expenditure went towards directly covering the cost of the regulation rather than innovation. Contrary to this conclusion, more recent work has supported that the "Weak" PH holds within the medium and long run. Rubashkina, Galeotti, and Verdolini present empirical evidence[8] from the European manufacturing sectors to show that pollution abatement and Control Expenditures induce innovation proxied by patent counts, consequently supporting the Weak PH. Albrizio, Kozluk, and Zipperer use OECD Environmental Policy Stringency data[9] and firm-level panel data to find that productivity gains often lag instead of occurring immediately. Within the EU ETS specifically, Martin, Muûls, de Preux, and Wagner[10] found limited evidence supporting significant productivity losses among regulated firms, with innovation effects remaining uneven. Calel and Dechezleprêtre[11] show that EU ETS exposure increased low-carbon patenting among regulated firms, supporting the weak PH over the long run. However, Abrell, Ndoye Faye, and Zachmann[12] argue that overallocation and low allowance prices significantly weakened compliance incentives, possibly limiting both productivity and inventive activity in the short run. Consequently, null effects may reflect weak policy implementation rather than a failure of Porter Hypothesis mechanisms. Thus, despite substantial scholarship on medium and long-run innovation effects, evidence on short-run responses to stringent environmental regulation is limited and mixed, especially during implementation that entails more immediate adjustment costs and compliance pressure.

This paper hypothesizes that stringent environmental regulation leads to no increase in productivity in regulated manufacturing sectors; rather, inventive activity increases, in line with the Weak Porter Hypothesis. The following discourse shall address the following. Section 2 will discuss the method and data used. Section 3 will provide the main results, while Section 4 will explore them in context. Section 5 will conclude.

Data

Data sources:

To examine the short-term effects of stringent environmental regulation on productivity within the European manufacturing sector, our study relied on a relatively large cross-country dataset constructed from the OECD Structural Analysis (STAN) database[13] and the EU Emissions Trading System (ETS) data viewer.[14] In addition to productivity outcomes, the effects on innovation activity were measured using the Eurostat Patent Application database.[15] To facilitate event studies or both the baseline productivity and innovation regressions, pre-2005 emissions data from the OECD Air Emissions Accounts Database[16] were used. The OECD STAN database is a tool for analyzing economic performance at the industry level, providing annual measures such as output, value added, and capital stock. In this paper, it was used to derive the country-year-industry panel of manufacturing activity, including metrics for Value added, Total Employment, and Labor productivity. The STAN sample used involves manufacturing industries from 23 EU countries observed annually over the period (1998-2010). These industries will be referred to by their International Standard Industry Classification (ISIC) Rev. 4 classification codes. Particular attention is paid to energy-intensive manufacturing industries, including Manufacture of paper and paper products (C17), Manufacture of coke and refined petroleum products (C19), Manufacture of other non-metallic mineral products (C23), and Manufacture of basic metals (C24). These manufacturing industries are among the most emissions-intensive and therefore plausibly have high ETS exposure, targeted by the introduction of the EU ETS in 2005. Now, with regard to the EU ETS, its introduction in 2005 represents a major tightening in environmental regulatory stringency. Specifically, a cap-and-trade system was introduced, requiring regulated industries to have allowances equal to their emissions. The sample extracted from the data viewer used verified emissions and allocated allowances for 23 EU countries at the Country-Year level, covering the period 2005-2010, corresponding to the early years of the ETS. The Eurostat Patent application database provided country-industry-year level data on the number of patent applications submitted each year. In this paper, data was used to derive data to proxy for innovation activity. The sample we extracted covered 1998-2010, spanning the same 23 EU countries and 4 manufacturing sectors as the STAN sample. We used the OECD Air Emissions Accounts database to provide pre-2005 emissions data, which was not available in the ETS database. Our sample contained country-industry-year level data for the 23 EU countries. The OECD STAN and EU ETS datasets were merged at the country-year level to create a final sample with stringency and productivity data. To ensure observations were subject to the same regulatory framework, OECD countries that were not present in the EU ETS data were excluded. In addition, some country-industry-year observations were dropped in cases where employment was reported as zero or value added as negative, thereby preventing the calculation of labor productivity. These restrictions mean that the analysis is based on economically meaningful data. Eurostat patent application data were also merged with EU ETS data at the country-year level to create a dataset with both patent and emissions data. The baseline productivity and innovation regressions are limited to 2005-2010 since the lagged EU ETS stringency proxy requires both verified emissions and allowance data, which are only reliably available post-2005. This results in 383 usable productivity observations and 375 usable innovation observations after exclusions. The event studies use a wider 1998-2010 panel and rely on pre-policy emissions exposure, with the productivity event study having 987 observations and the innovation event study having 975 observations.

Variable Construction:

The key independent variable was environmental regulation stringency. Direct sector-country-year measures of environmental regulatory stringency are particularly limited, especially within the EU ETS framework. Thus, effective environmental regulation stringency has been proxied with the ratio of verified emissions to allocated allowances under the EU ETS. Stringency is first constructed as the emissions-to-allowances ratio lagged by one year, then log-transformed for the regression. The measure used captures the effective bindingness of emissions regulation, since sectors with verified emissions that are close to or exceeding allocated allowances face greater compliance pressure, and thereby stronger incentives to improve efficiency.

$$LogStringency_{ic,t-1} = log(\frac{VerifiedEmissions_{ic,t-1}}{AllocatedAllowances_{ic,t-1}})$$

While the proxy used reflects the realized compliance pressure instead of legislative strictness, it more clearly captures regulatory burden within the EU ETS framework than wider environmental policy datasets such as the OECD Environmental Policy Stringency (EPS) Index.[17] OECD EPS measures national-level policy settings, but lacks the sector-country-year granularity and does not reflect sector-specific compliance pressure faced by regulated industries. By using the ratio of verified emissions to allocated allowances, the degree to which regulated sectors approach exceeding emissions caps is directly reflected, therefore serving as a more precise indicator of compliance pressure under cap-and-trade regulation. That said, it may partially reflect differences in firm efficiency or emissions intensity. As such, alternative measures such as carbon pricing indicators and wider environmental policy stringency indices were considered, but were not used as they lacked sufficient industry-country-year granularity for this analysis. To mitigate simultaneity concerns, the variable is logged and lagged by one year, though this does not fully eliminate possible endogeneity. Results should be interpreted as better reflecting effective regulatory pressure rather than pure legislative stringency. To assess the sensitivity of findings to variable construction and reduce concerns regarding possible outlier distortion, robustness checks were performed using a winsorized variation of the stringency ratio as an alternative specification of environmental regulation stringency.

The dependent variable used to test productivity in this study was labor productivity, measured as the logarithm of value added per worker. This is widely used in Porter Hypothesis literature as an indicator of efficiency. Specifically, productivity is constructed using value added (VA) and total employment from the OECD STAN database as

$$Productivity_{ict} = log(VA_{ict}) - log(employment_{ict})$$

By using logarithms, can reduce the skewness in the data while also making productivity levels comparable across heterogeneous units.

To proxy innovation, we used patent applications as a measure of inventive activity and technological change. This allows for a broad and observable proxy for new inventive output, though it is unable to capture factors such as innovation quality or non-patented innovation. The specific construction of the proxy is below.

$$Innovation_{ict} = log(Patent\ Applications_{ict} + 1)$$

By adding +1, we avoided any invalid results from log(0).

Country-industry fixed effects control for time-invariant structural differences across sectors, and year fixed effects capture common macroeconomic shocks over time.

Summary Statistics:

Table 1: Summary statistics for productivity regression. The statistics indicate substantial variation in productivity across countries and industries. The negative mean value of lagged log stringency suggests that emissions were generally below allocated allowances during the sample period.

Summary Statistics: Productivity Regression Sample					
	Mean	SD	Min	Max	N
productivity	4.86	1.86	-1.63	11.84	383
log_stringency_lag1	-0.19	0.18	-0.96	0.27	383
log_employment	2.58	1.77	-5.81	5.64	383

Table 1 reports the summary statistics for the main variables used in the empirical analysis. Labor productivity has a mean of 4.86 (SD = 1.86), indicating substantial variation across countries, industries, and time. Lagged log stringency, defined as the log of verified emissions divided by allocated allowances (t-1), had a mean of -0.19 (SD = 0.18), indicating that emissions were often below allocated allowances. This suggests that in the early ETS period, many industries faced relatively weak compliance pressure.

Log employment has a mean of 2.58 (SD = 1.77), suggesting heterogeneity in industry size. The sample consists of 383 observations over the period of 2005-2010.

Table 2: Summary statistics for innovation regression. Patenting activity varies considerably across observations. The relatively low average stringency value suggests that many industries faced limited compliance pressure during the early ETS period.

Summary Statistics: Innovation Regression Sample					
	Mean	SD	Min	Max	N
log_patents	1.82	1.54	0.00	5.97	375
log_stringency_lag1	-0.19	0.18	-0.96	0.27	375

Table 2 shows the summary statistics for the variables used in the innovation regression. The mean value of log patents is 1.82 (SD = 1.54), indicating substantial variation in patent activity. Lagged log stringency had a mean of -0.19 (SD = 0.18), indicating that emissions were often below allocated allowances. The sample consists of 375 observations over the period of 2005-2010. The smaller sample size relative to the productivity regressions was caused by missing patent observations in some country-industry-year groups.

■ Empirical Strategy

Baseline Two-Way Fixed Effects Specification:

To estimate the short-run relationship between stricter environmental regulation and productivity within the manufacturing sector, this paper employs a panel fixed effects style framework with two-way fixed effects and country-industry linear time trends.

DOI: 10.36838/IJHSR820.8

$$Productivity_{ict}=\beta log(stringency_{c,t-1}) +\gamma log(Employment_{ict}) \alpha_{ic}+\lambda_t \delta_{ic} t+\varepsilon_{ict}$$

Where $Productivity_{ict}$ represents the labor productivity in industry i, country c, and year. $log(stringency_{,c,t-1})$ denotes lagged log stringency, measured as the logarithm of verified emissions divided by allocated allowances. $Employment_{ict}$ controls for employment. Industry-country fixed effects (α_{ict}) absorb differences in technology, structure, and productivity levels across country-industry pairs that do not change over time. Year fixed effects (λ_t) absorb macro shocks that affect all manufacturing sectors and countries. Country-industry specific linear time trends ($\delta_{ic}t$) account for differential pretrends across country-industry pairs. Lastly, the error term (ε_{ict}) accounts for unobserved factors affecting productivity that vary across country, industry, and time. Identification comes from within-country-industry variation over time. The final panel includes 58 country-industry units observed across 3 post-treatment years.

Identification and Interpretation:

The two-way fixed effects used in the model help to isolate the effect of increasing the $log(stringency_{,c,t-1})$ of environmental regulation by removing confounding factors that are constant over time within country-industry pairs. The coefficient of interest in this case is β, as it captures the short-run response of productivity to the EU ETS Scheme.

The approach used in this study provides quasi-causal experimental evidence under standard panel data assumptions, although results should be interpreted as short-run associations instead of purely causal effects. Accordingly, the results provide insight into short-run industry-level adjustment and compliance responses to the 2005 EU ETS regulation, rather than long-run innovation-driven impacts.

Event Study design:

To improve the credibility of the empirical model, the initial analysis for both productivity and innovation was complemented with an event study, focusing on the introduction of the EU ETS in 2005. In particular, the event study tests the parallel trends assumption and provides insight into the dynamics of productivity before and after 2005. The specification is given by

$$Productivity_{ict}= \sum_{k\neq -1} \beta_k (average\ emissions_{1998-2004} x1[event\_time = k]) + \alpha_{ic}+\lambda_t+\delta_{ic} t+\varepsilon_{ict}$$

Where $event\_time = t$-2005 indexes year relative to the introduction of the EU ETS. $event\_time = k$ is a dummy variable that equals one if year t is k years away from 2005. ETS exposure captures time-invariant differences in regulatory exposure, measured using average emissions intensity in the pre-policy period (1998–2004). The pre-treatment year ($k-1$) is removed from the regression as it acts as the baseline.

A separate coefficient β_k is estimated for each year relative to the introduction of the ETS. The coefficients for pre 2005 years test for differential trends before the policy shock.

Innovation empirical strategy:

A second two-way fixed effects baseline model was estimated to examine whether the EU ETS induced innovative activity in the short run. To do this, the baseline specification was re-estimated with innovation as the dependent variable:

$$Innovation_{ict}=\beta\ logStringency_{c,t-1} + \alpha_{ic}+\lambda_t+\delta_{ic} t+\varepsilon_{ict}$$

Where $innovation_{ict}$ represents the level of innovation in industry i , country c , and year. As in the productivity model, $log(stringency_{,c,t-1})$ is the measure of environmental regulation stringency, proxied by lagged log regulatory stringency. Furthermore, industry-country, year, and country-industry fixed effects are used in the same manner as the productivity baseline regression. The error term accounts for unobserved factors.

To assess the Weak Porter Hypothesis, the same event study framework was applied using innovation as the dependent variable.

$$Innovation_{ict}=\sum_k \beta_k (average\ emissions_{1998-2004} x1[event = k]) + \alpha_{ic}+\lambda_t+\delta_{ic} t+\varepsilon_{ict}$$

This innovation event study examines whether more highly exposed manufacturing sectors had pre-trends
in inventive activity before the introduction of ETS in 2005.

Limitations:

The chosen data and methodology used have a couple of important limitations. The use of country-industry and year fixed effects improves validity by controlling for unobserved heterogeneity, but combined with the relatively limited sample size (382 observations), may reduce statistical precision. Therefore, standard errors may be higher, limiting the study's ability to detect smaller short-run effects. In addition, the proxy used reflects realized compliance pressure, instead of legislative strictness. While it is practical, given data constraints, this measure may also partially capture differences in efficiency. Similarly, patent applications offer an observable measure of inventive activity, but may not fully capture innovation quality, process improvements, or non-patented technological adaptations. Moreover, because this study focuses on short-run effects within EU ETS-regulated sectors, findings may not generalize to long-run outcomes or non-European policy environments.

Results

Our baseline model estimated the association between labor productivity and environmental regulation stringency, proxied by lagged log stringency.

In Table 3, we report the results from the baseline and robustness regression using a two-way fixed effects specification with country-industry fixed effects. Column (1) reports the baseline estimate where the coefficient on lagged regulatory stringency is 0.005 (SE = 0.187), indicating no short-run relationship between lagged log stringency and labor productivity. In Column (2), adding country-industry linear trends increases the coefficient to 0.011(SE = 0.306), but it remains statistically insignificant. Columns (3) and (4) test robustness by restricting the time period to phase 1 of ETS (2005-2007) and excluding the crisis years 2008-2009. The estimated effects (0.052 and 0.025, respectively) remain statistically insignificant. Across all specifications, Log employment is negative and statistically significant, which reveals that within a country-industry, years with more workers are associated with lower productivity. This is consistent with diminishing returns. Overall, the baseline re-

sults suggest that stricter environmental regulation does not result in short-run productivity gains.

Table 3: Two-way fixed effects regressions of labor productivity. Across all specifications, the estimated effect of regulatory stringency remains small and statistically insignificant. These results provide little evidence that stricter environmental regulation increased short-run productivity.

Productivity: Baseline and Robustness Checks

	Baseline	+ IC trends	Phase I (05–07)	Drop 08–09
log_stringency_lag1	0.005	0.011	0.052	0.025
	(0.204)	(0.306)	(0.331)	(0.147)
log_employment	-1.016***	-0.745***	-3.497*	-1.253*
	(0.174)	(0.084)	(1.591)	(0.506)
Num.Obs.	383	383	154	229
R2	0.972	0.984	0.985	0.977
R2 Adj.	0.965	0.973	0.969	0.965
R2 Within	0.167	0.092	0.308	0.195
R2 Within Adj.	0.161	0.084	0.289	0.184
AIC	353.7	297.0	154.6	232.1
BIC	681.4	928.7	397.6	510.2
RMSE	0.31	0.23	0.24	0.28
Std.Errors	by: ref_area	by: ref_area	by: ref_area	by: ref_area
FE: ref_area^activity	X	X	X	X
FE: year	X	X	X	X
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001				
Country-industry fixed effects: 58				
Year fixed effects: 3 years				
Standard errors clustered at the country-industry level (58 clusters)				

Each column contains estimates from regressions with country–industry and year fixed effects. Standard errors (SE) are clustered at the country level and reported in parentheses.

Graphical evidence:

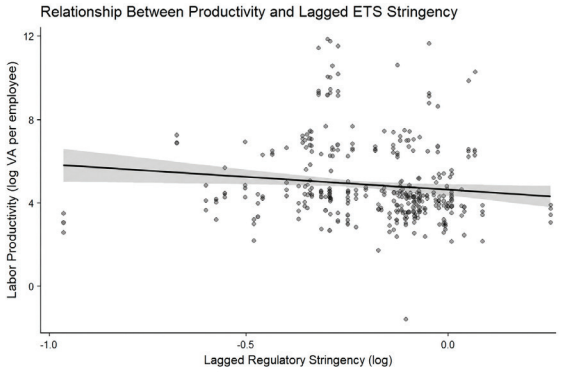


Figure 1: Productivity vs lagged regulatory stringency. The fitted relationship is slightly negative and weak. This suggests that cross-sectional differences in regulatory stringency alone do not strongly predict productivity outcomes.

Figure 1 plots labor productivity against lagged regulatory stringency and creates a slightly negative fitted relationship. Furthermore, the scatter plot revealed substantial dispersion at similar emission levels, suggesting that the cross-sectional variation alone did not tightly predict productivity when comparing different country-industry pairs.

Figure 2 shows lagged regulatory stringency over the period 2006-2010. Emissions followed a modest upward trend initially before sharply declining around 2008-2009, coinciding with both the tightening of ETS rules and the global financial crisis. This graph highlights the meaningful changes in emissions intensity. However, this paper focuses largely on the short-term effects, reflecting the period 2005-2006.

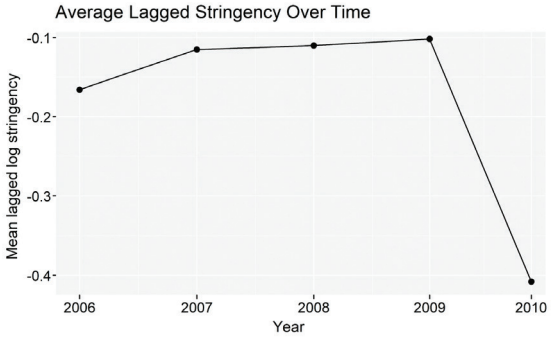


Figure 2: Average ETS emissions over time (2005-2010). Average ETS emissions increased slightly before declining sharply after 2008. This indicates a substantial change in emissions levels during the latter part of the sample period.

Even so, fixed effects absorb common macroeconomic shocks.

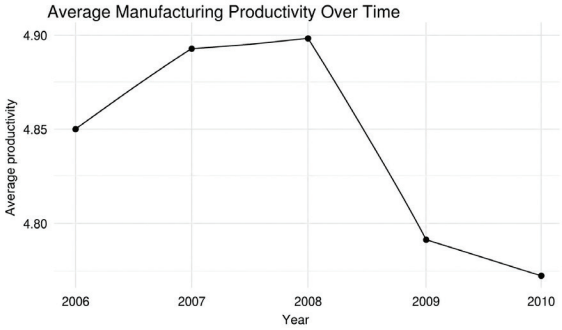


Figure 3: Average manufacturing productivity over time (2005-2010). Productivity increased slightly between 2006 and 2008 before declining noticeably after 2008. Overall, there is no clear evidence of sustained productivity growth following the introduction of the EU ETS.

Event-study Evidence:

To examine whether productivity trends differed by ETS exposure before the policy shock and to examine short-run dynamics after the start of the ETS, an event-study specification was run using panel data from 1998 to 2010.

Table 4 reports the full set of event-time coefficients from this specification, including standard errors clustered at the country level and controlling for country-industry and year fixed effects. Across event times -7 to -2, all coefficients were close to zero and thereby statistically insignificant. This means that there is no evidence of differential pre-trends in productivity across ETS-exposed industries before 2005. Furthermore, around the policy implementation and in subsequent years, the coefficients remain close to zero, suggesting that higher-pre ETS exposure is not associated with significant short-run changes in productivity following the introduction of ETS. In short, the event study supports the baseline findings, showing no significant changes to short-run productivity after the introduction of ETS.

DOI: 10.36838/IJHSR820.8

Table 4: Event-Study regression results. Both pre-treatment and post-treatment coefficients remain close to zero and statistically insignificant. This suggests no detectable short-run productivity response to ETS exposure.

Event Study Results: Productivity and Pre-ETS Exposure	
Event study (Productivity)	
event_time = -7 × Exposure_ic	2.3e-06 (1.8e-06)
event_time = -6 × Exposure_ic	2.5e-06 (1.9e-06)
event_time = -5 × Exposure_ic	2.5e-06 (1.4e-06)
event_time = -4 × Exposure_ic	2.2e-06 (1.2e-06)
event_time = -3 × Exposure_ic	1.6e-06 (1.2e-06)
event_time = -2 × Exposure_ic	8.7e-07 (1.0e-06)
event_time = 0 × Exposure_ic	-2.1e-07 (1.0e-06)
event_time = 1 × Exposure_ic	6.1e-07 (1.3e-06)
event_time = 2 × Exposure_ic	-7.8e-07 (1.2e-06)
event_time = 3 × Exposure_ic	-1.8e-06 (1.3e-06)
event_time = 4 × Exposure_ic	-1.7e-06 (1.1e-06)
event_time = 5 × Exposure_ic	-1.8e-06 (1.3e-06)
Num. Obs.	987
R2	0.959
R2 Within	0.044
Std. Errors	Clustered by: ref_area
FE: REF_AREA^ACTIVITY	Yes
FE: YEAR	Yes
Country-industry fixed effects: 58	
Year fixed effects: 3	
Standard errors clustered at the country-industry level (58 clusters)	

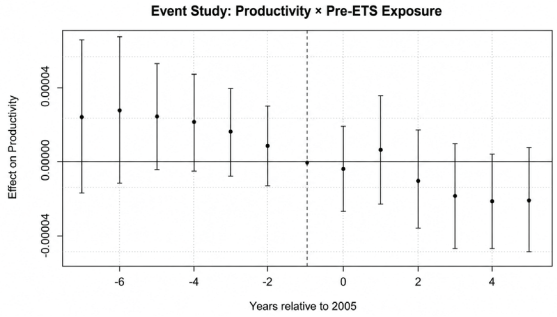


Figure 4: Event study plot. The estimated coefficients fluctuate around zero, and their confidence intervals generally include zero. This provides visual support for the absence of significant productivity effects following ETS implementation.

Figure 4 plots the estimated event-study coefficients consistent with Table 4.

Innovation results:

Table 5 reports the baseline and robustness estimates of the relationship between lagged log stringency and innovation, proxied by log patents. The coefficient on lagged log stringency was 0.045 with a SE 0.170) and statistically insignificant, which implies that there was no short-term increase in innovation activity. By adding country-industry linear trends in Column (2), we found a negative but imprecise coefficient of -0.325 (SE = 0.173), and restricting estimates to 2005-2007 provided a similar estimate of - 0.233. All coefficients are statistically insignificant and indistinguishable from zero, suggesting that increased lagged log stringency did not lead to any meaningful increase in short-run patenting.

Table 5: Innovation baseline regression results. The estimated effects of regulatory stringency on patenting are statistically insignificant across all specifications. These findings provide limited support for the Weak Porter Hypothesis in the short run.

Innovation: Baseline and Robustness Checks				
	Baseline	+ IC trends	Phase I (05–07)	Drop 08–09
log_stringency_lag1	0.045	-0.325+	-0.233	0.217
	(0.170)	(0.173)	(0.318)	(0.241)
Num.Obs.	375	375	150	225
R2	0.967	0.981	0.983	0.969
R2 Adj.	0.959	0.967	0.965	0.953
R2 Within	0.000	0.012	0.003	0.008
R2 Within Adj.	-0.003	0.007	-0.010	0.001
AIC	261.3	217.3	97.8	202.5
BIC	575.5	826.0	329.6	468.9
RMSE	0.28	0.21	0.20	0.27
Std.Errors	by: ref_area	by: ref_area	by: ref_area	by: ref_area
FE: ref_area^activity	X	X	X	X
FE: year	X	X	X	X
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001				
Country-industry fixed effects: 58				
Year fixed effects: 3 years				
Standard errors clustered at the country-industry level (58 clusters)				

Table 6: Innovation event study regression results. Pre-policy coefficients show no evidence of differential trends, while post-policy coefficients remain close to zero. This suggests that ETS exposure did not generate measurable short-run increases in patenting activity.

Event Study Results: Innovation and Pre-ETS Exposure	
Event study (Innovation)	
event_time = -7 × Exposure_ic	-1.1e-06 (1.2e-06)
event_time = -6 × Exposure_ic	9.5e-07 (1.0e-06)
event_time = -5 × Exposure_ic	2.2e-07 (1.1e-06)
event_time = -4 × Exposure_ic	8.3e-07 (1.0e-06)
event_time = -3 × Exposure_ic	5.4e-07 (1.0e-06)
event_time = -2 × Exposure_ic	1.2e-06 (1.0e-06)
event_time = 0 × Exposure_ic	9.9e-07 (1.0e-06)
event_time = 1 × Exposure_ic	1.3e-06 (1.0e-06)
event_time = 2 × Exposure_ic	9.1e-07 (1.0e-06)
event_time = 3 × Exposure_ic	7.1e-07 (1.1e-06)
event_time = 4 × Exposure_ic	5.5e-07 (1.1e-06)
event_time = 5 × Exposure_ic	7.1e-07 (1.2e-06)
Num. Obs.	975
R2	0.961
R2 Within	0.011
Std. Errors	Clustered by: ref_area
FE: REF_AREA^ACTIVITY	Yes
FE: YEAR	Yes
Country-industry fixed effects: 58	
Year fixed effects: 3	
Standard errors clustered at the country-industry level (58 clusters)	

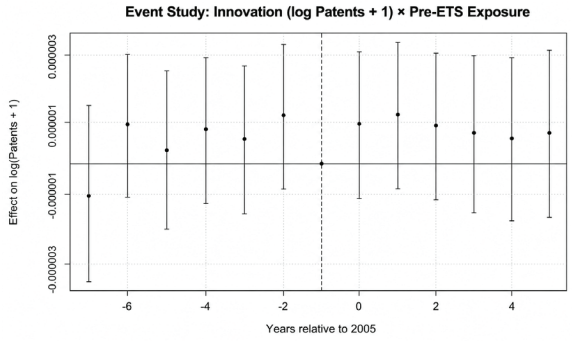


Figure 5: Innovation event study regression plot. The estimated effects remain small and statistically insignificant throughout the study period. This visual evidence is consistent with the baseline innovation regressions.

In order to confirm our conclusion, we ran an event study to examine pre-trends. Table 6 and Figure 5 present the results. The coefficients in the pre-policy period (event times -7 to -2) were small in magnitude and statistically insignificant, thereby providing no evidence of differential pre-trends. Following the introduction of ETS, coefficients remain close to zero and statistically insignificant across all post-policy event times. Overall, the results are consistent with, but do not prove, the parallel trends assumption, thereby supporting the baseline regression results. These results provide very limited evidence in support of the Weak Porter's hypothesis within the short run, as innovative activity did not experience a meaningful increase.

Robustness Checks:

To assess the stability of the baseline findings, a lagged winsorized logged specification of environmental regulation stringency was used, while retaining the same fixed-effects and standard-error framework. The construction of the variable can be seen below:

$$WinsorisedLogStringency_{ic,t-1} = lag(Winsorise(log(\frac{VerifiedEmissions_{ic,t-1}}{AllocatedAllowances_{ic,t-1}})))$$

As shown in Table 7, the coefficient on lagged winsorized log stringency was negative for productivity (-0.019) and positive for innovation (0.046), yet both remained statistically insignificant. Therefore, the winsorized logged specification produced results that were consistent with the baseline regressions. This suggests that the baseline findings are robust to alternative treatment of extreme observations and are not disproportionately driven by outlier values.

Table 7: Winsorized robustness check results. The coefficients remain statistically insignificant after reducing the influence of extreme observations. This suggests that the baseline findings are not driven by outliers.

Robustness Check: Winsorized Logged Environmental Stringency

	Productivity Robustness	Innovation Robustness
Lagged Winsorized Log Stringency	-0.019	0.046
	(0.015)	(0.043)
Observations	383	383
R^2	0.996	0.984
Adjusted R^2	0.996	0.981
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001		
Country-industry fixed effects: 58		
Year fixed effects: 3 years		
Standard errors clustered at the country-industry level (58 clusters)		

In addition to this, to assess whether null findings were driven by limited statistical power, simplified model specifications were run with fewer controls. Only country-industry and year fixed effects were run with clustered SEs.

Table 8: Simplified model specification results. The simplified specifications produce estimates similar to the main models. This reinforces the conclusion that regulatory stringency had no detectable short-run effect on productivity or innovation.

Robustness Check: Simplified Model Specification

	Productivity Simplified	Innovation Simplified
Lagged Logged Stringency	-0.004	0.008
	(0.003)	(0.010)
Observations	383	383
R^2	0.996	0.984
Adjusted R^2	0.996	0.980
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001		
Country-industry fixed effects: 58		
Year fixed effects: 3 years		
Standard errors clustered at the country-industry level (58 clusters)		

As observed in Table 8, the coefficient on productivity was -0.004, and the coefficient on innovation was 0.008, with both again being statistically insignificant. This supports the baseline null findings.

Discussion

The primary goal of this paper was to explore the relationship between environmental regulation stringency and productivity within the short run. We hypothesized that more stringent regulation would not immediately lead to increases in productivity; rather, inventive activity would experience an increase. Our study combined panel data from 22 OECD countries to show that, within the short run, more stringent regulation was not statistically associated with either increased productivity or innovation. Accordingly, the results do not support either the Strong or Weak Porter's Hypothesis in the short run. The baseline productivity regression results provided no evidence that lagged log stringency had a statistically significant effect on productivity once employment was controlled for. With regard to the event study, the pre-policy coefficients are close to zero and statistically insignificant, providing little evidence of differential pretrends. This strengthens the credibility of the empirical design and supports the interpretation of the baseline results as short-run associations. This finding is consistent with previous literature that suggests compliance pressure is often major, coupled with limited efficiency gains. Notably, Greenstone showed that manufacturing activity often decreases18 immediately after a policy shock, consistent with short-run adjustment costs. Overall, the absence of short-run productivity gains was expected and matched the consensus of previous research regarding the strong PH. To determine whether the absence of productivity gains reflects the time required for new technologies to materialize or if, instead, firms were focused on cost abatement, this paper also examined patent applications as a proxy for innovation. Our baseline regression showed no significant increase in inventive activity following the introduction of ETS. The innovation event study similarly supported the parallel trends assumption, as coefficients were practically indistinguishable from zero, in-

DOI: 10.36838/IJHSR820.8

dicating no evidence of differential pre-trends. This suggests that the identifying assumptions are reasonably satisfied, although our estimates should still be interpreted as short-run associations rather than causal effects. While limited productivity gains were expected, the lack of an increase in inventive activity was surprising. One plausible explanation is that Phase 1 of the EU ETS may have generated insufficient compliance pressure as a result of generous allowance allocations, thereby weakening the incentives for immediate innovation. This interpretation is in line with Abrell *et al.*[19] and explains the divergence between our short-run null findings and Calel & Dechezleprêtre's medium-run findings. As such, our results contrast with some other studies that test for innovation. One such paper, Calel and Dechezleprêtre, observed an increase in low-carbon patenting.[9] among firms exposed to ETS over the medium run, supporting the Weak PH. The results suggest that even if regulation encourages firms to innovate, inventive activity may not occur immediately, with possible productivity gains emerging even later. Based on previous research focusing on the medium and long run, Innovation frequently takes time to translate into output efficiency due to adjustment costs. Our study has relevant implications for the field. Literature in this field tends to focus heavily on the medium and long run, leaving holes in understanding within the short-run. Importantly, this paper helps to bridge the gap in the literature by revealing that any productivity increases in line with the Strong PH do not occur within the short run. Literature in this field tends to agree that the Weak Porter's hypothesis, wherein Innovation increases in response to regulation, is valid. Yet, we found that more stringent regulation has little impact on inventive activity proxied by patents. Therefore, we can reasonably assume that inventive activity did not increase right after the introduction of ETS. By explicitly focusing on the short-run productivity and innovation effects, the results suggest that gains are not automatic. This has important implications for policy-makers, as our results suggest that expecting quick productivity gains is unreasonable. Instead, environmental policies may take longer to become effective. Policymakers should also avoid assuming that the effects of environmental regulation are homogeneous and will surely generate innovation or productivity gains in the short-run. Our results suggest that policymakers should look to make long-term policy commitments and sufficiently binding compliance in order to see a significant increase in technological improvements within industry. However, some limitations limit the interpretation of our results. First, while both event study estimates show little evidence of significant pre-trends, the relatively short time window constrains causal interpretation. Therefore, estimates should be interpreted cautiously as short-run associations rather than definitive causal effects. Second, patent applications might not fully capture changes in inventive activity. Future work in this field could explore the heterogeneity of policies and their effect, examine short-run effects from a firm-level perspective, or focus on technology spillovers. Lastly, given the relatively limited post-treatment sample and the extensive fixed-effects structure, statistical power may be constrained, so null findings should be interpreted cautiously. In all, neither the Weak nor the Strong Porter's Hypothesis holds in the short run.

Conclusion

In conclusion, this paper finds no evidence in support of either the Strong or Weak Porter's Hypothesis within the short run. Our estimations find that increased lagged log stringency is not statistically associated with higher productivity or innovation over the 2005-2007 period. Furthermore, the event studies show little evidence of differential pre-trends, thereby supporting the parallel trends assumption. However, it is worth noting that the short time window limits causal interpretation. As such, all our estimates should be viewed as descriptive short-run associations, rather than clear causal effects. These findings challenge the assumption introduced at the beginning of this paper that environmental regulation necessarily creates innovation and further benefits as a result. Instead, firms facing such constraints likely prioritize cost-adjustment, which delays the realization of the Porter-style gains. With regard to policy implications, our evidence suggests that immediate productivity or innovation gains should not be expected shortly following the introduction of environmental regulation. Moreover, in the short run, such regulation may be insufficient to rapidly stimulate innovation unless paired with measures such as transitional subsidies. Future research should therefore focus on the regulatory conditions under which Porter-style gains may arise in the short run.

Acknowledgments

I want to express gratitude to my mentor, Dr. Ceyhun Elgin from Columbia University, for his guidance and detailed feedback during the research process. I also want to thank Daryn Howland for her support.

References

1. Schumpeter, J. A. Capitalism, Socialism and Democracy; Harper & Brothers: New York, 1942.
2. Romer, P. M. Endogenous Technological Change. *J. Polit. Econ.* **1990**, *98* (5), S71–S102.
3. Porter, M. E.; van der Linde, C. Toward a New Conception of the Environment–Competitiveness Relationship. *J. Econ. Perspect.* **1995**, *9* (4), 97–118. https://doi.org/10.1257/jep.9.4.97.
4. Jaffe, A. B.; Newell, R. G.; Stavins, R. N. Technological Change and the Environment. In *Handbook of Environmental Economics*; Mäler, K.-G., Vincent, J. R., Eds.; Elsevier: Amsterdam, 2003; Vol. 1; pp 461–516.
5. Gray, W. B.; Shadbegian, R. J. Plant Vintage, Technology, and Environmental Regulation. *J. Environ. Econ. Manage.* **2003**, *46* (3), 384–402. https://doi.org/10.1016/S0095-0696(03)00031-7.
6. Rubashkina, Y.; Galeotti, M.; Verdolini, E. Environmental Regulation and Competitiveness: Empirical Evidence on the Porter Hypothesis from European Manufacturing Sectors. *Energy Policy* **2015**, *83*, 288–300. https://doi.org/10.1016/j.enpol.2015.02.014.
7. Albrizio, S.; Kozluk, T.; Zipperer, V. Environmental Policies and Productivity Growth: Evidence across Industries and Firms. *J. Environ. Econ. Manage.* **2017**, *81*, 209–226. https://doi.org/10.1016/j.jeem.2016.06.002.

8. Martin, R.; Muûls, M.; de Preux, L. B.; Wagner, U. J. Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme. *Am. Econ. Rev.* **2014**, *104* (8), 2482–2508. https://doi.org/10.1257/aer.104.8.2482.
9. Calel, R.; Dechezleprêtre, A. Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Rev. Econ. Stat.* **2016**, *98* (1), 173–191. https://doi.org/10.1162/REST_a_00470.
10. Abrell, J.; Ndoye Faye, A.; Zachmann, G. Assessing the Impact of the EU ETS Using Firm-Level Data; Bruegel Working Paper No. 2011/08; Bruegel: Brussels, 2011.
11. Dechezleprêtre, A.; Sato, M. The Impacts of Environmental Regulations on Competitiveness. *Rev. Environ. Econ. Policy* **2017**, *11* (2), 183–206. https://doi.org/10.1093/reep/rex013.
12. Ambec, S.; Cohen, M. A.; Elgie, S.; Lanoie, P. The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Rev. Environ. Econ. Policy* **2013**, *7* (1), 2–22. https://doi.org/10.1093/reep/res016.
13. Organization for Economic Co-operation and Development (OECD). *Structural Analysis (STAN) Database*; OECD Publishing: Paris, France.
14. European Commission. *European Union Emissions Trading System (EU ETS) Data Viewer*; Directorate-General for Climate Action: Brussels, Belgium.
15. Eurostat. *Patent Applications to the European Patent Office (EPO) by Priority Year*; European Commission: Luxembourg.
16. Organization for Economic Co-operation and Development (OECD). *Air Emissions Accounts*; OECD Publishing: Paris, France.
17. Organization for Economic Co-operation and Development (OECD). *Environmental Policy Stringency Index (EPS)*; OECD Publishing: Paris, France.
18. Greenstone, M. The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments. *J. Polit. Econ.* **2002**, *110* (6), 1175–1219.
19. Hall, B. H.; Lotti, F.; Mairesse, J. *Innovation and Productivity in SMEs: Empirical Evidence for Italy*; NBER Working Paper No. 17329; National Bureau of Economic Research: Cambridge, MA, 2011.

Authors

Chris Kumala is a student at Rangitoto College with a strong interest in the intersection of economics and sustainability. His academic focus centers on endogenous growth theory and the circular economy, particularly how economic systems can drive long-term innovation while supporting environmental sustainability.

Appendix

Appendix A: Sample Construction

Step	Productivity	Innovation
Raw observations	1078	1078
EU ETS merge	1001	1001
Missing employment/value added exclusions	987	987
Patent missing exclusions	987	975
Lag losses	858	846
Final TWFE	383	375
Final Event Study	987	975

DOI: 10.36838/IJHSR820.8

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

How Can We Build an AI To Diagnose Textually Ambiguous Mental Health Disorders from Subjective Patient Reports?

Dhairya Jain

Jayshree Periwal Global School, Jagatpura, Jaipur, Rajasthan 302033, India; jaindhairya21@gmail.com

ABSTRACT: Nearly 1 in 7 people in the world live with a mental disorder. Still, in many regions, only 7% of individuals get help due to a big treatment gap due to a lack of trained health professionals, cost issues, or stigma, especially in low- and middle-income countries. However, the advancements in the field of AI have been a big reason for the rise in interest in AI-based mental health screening tools that can analyze patient-generated texts, such as clinical notes, questionnaires, and social media posts. Previous studies in this domain have consistently demonstrated that AI systems often achieve strong results when diagnosing well-defined mental disorders such as major depressive disorder and generalized anxiety disorder, because they exhibit consistent linguistic markers and standard patient-reported outcomes. However, comparatively less progress has been made for mental health disorders characterized by overlapping and ambiguous symptoms, such as PTSD and subtypes of depression. These challenges highlight a critical limitation in current AI approaches, particularly in handling linguistic ambiguity and subjective patient-reported data. This paper presents a secondary analytical review of existing literature, focusing on model architectures, data sources, and evaluation frameworks, and proposes an uncertainty-aware AI architecture incorporating ambiguity detection, multi-label prediction, and human-centered evaluation.

KEYWORDS: Robotics and Intelligent Machines, Machine Learning, Mental Health, Diagnosis.

Introduction

The use of Artificial Intelligence (AI) in the mental health field has been growing for the last few years, especially in the areas of machine learning and natural language processing.[1-3] An AI system, through analyzing different forms of written language (clinical notes, self-report questionnaires, online text),[4] has shown that it can detect mental health "signals" and come close to matching the performance of clinicians when diagnosing some well-defined mental health conditions (major depressive disorder, generalized anxiety disorder, etc.)[5]

These studies indicate that it may be possible to leverage AI tools to help expand access to mental health services on a larger scale through the use of Text-Based AI. Certain mental disorders can be easily classified or diagnosed using an AI tool. On the other hand, many disorders are difficult to classify or diagnose. The language used to describe depression subtypes, PTSD, etc, is very subjective, with a great deal of overlap between types of disorders, and will vary greatly between individuals. Because the majority of the information for diagnosing these disorders will come from patient-reported outcomes, and because these disorders do not have any objective diagnostic markers, they are often not diagnosed until the patient presents later in the course of a clinical presentation.[6] Therefore, most of the current research on AI to date has treated the diagnosis of these disorders as a relatively straightforward classification task, placing little to no emphasis on factors like uncertainty, the need for explanation, and the fact that the majority of self-reported mental health experiences are very subjective. We still don't have a good understanding of the ways in which AI systems should be developed to consider the ambiguity of language and how to more effectively incorporate Patient Reported Outcomes (PROs) into the development of AI systems.

Building AI systems that can diagnose and treat mental health disorders using inherently vague or ambiguous text will require examining three areas of research: Model Architecture, Model Training, and Model Evaluation. Recent advancements in the field of AI, such as chatbots, have shown the revolutionary potential of AI in the mental health field. While many chatbots are accurate in diagnosing disorders based on ideal datasets, their accuracy often decreases when used in the real world with subjective datasets. There is a lot of scope in filling this gap, as nowadays the demand for healthcare services is ever-increasing, and many countries are experiencing a shortage of healthcare practitioners, especially physicians.[7] This shortage is very crucial to consider, as this often leads to delayed treatments, misclassifications, and poor outcomes. Addressing this issue can help us to improve early identification of disorders, and further in making a responsible development of an AI specific to mental health contexts. Currently, we have only one form of input, and that is text, but in the future, we can add speech, behavior, and facial expressions as part of the input. This study carries out a secondary analysis of existing AI methods in mental health research. Previous studies are organized by disorder type, data source, and evaluation metrics. The model will be evaluated based on both traditional machine learning measures and human-based evaluation measures.

Traditional measures of performance include accuracy, precision, recall, and the F1 -score, which assess the statistical accuracy of disorder classification as observed in other literature. Besides, human-centered indicators will be included through a Likert scale (an ordered rating system used in

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.9

surveys, Strongly agree - strongly disagree) to measure judgments of clarity, usefulness, appropriateness, and reliability, thus summarizing how patients and clinicians perceive the quality and credibility of the AI-produced screening results.[8] This study adopts a secondary analytical review methodology, synthesizing existing literature in AI-based mental health diagnostics. Sources were selected based on relevance to model architectures, data inputs, and evaluation frameworks, with an emphasis on peer-reviewed academic research and established institutional reports.

Discussion

This paper has demonstrated how artificial intelligence works well in spotting mental health issues like major depression and anxiety, where speech patterns and symptoms are clear and consistent. Yet performance drops when dealing with conditions marked by unclear symptoms or overlapping traits such as PTSD or less typical forms of depression, where language may be subtle or indirect. Previous research papers tend to focus too much on numbers like accuracy rates instead of deeper concerns like how uncertain the model might be. Interpretation matters just as much, yet it often takes a back seat in current work. Even when models score high technically, they do not always earn user confidence if explanations feel opaque or irrelevant. Trust builds differently among individuals using chatbots or online platforms meant for emotional support. Clarity of response and whether people find them helpful shape actual usage far more than backend precision ever could. Outcomes depend heavily on personal experience, something algorithms rarely capture fully. Moving forward, intelligent systems should assist professionals rather than replace them. Including self-reports from patients alongside transparent results can guide better decisions. Oversight by trained clinicians ensures risks stay low while keeping care grounded in real needs.

Mental health diagnosis

How mental health disorders are currently diagnosed:

Mental health disorders mainly rely on clinical interviews and patient-reported outcomes, rather than visible symptoms like physical disorders. To diagnose a particular mental health disorder, clinicians use various diagnostic frameworks, such as DSM-5 and ICD-11,[9] which are a few examples of these frameworks, as they assess symptoms based on how severe and how long a symptom has been present. Various questionnaires, such as PHQ-9 and GAD-7, also help clinicians in diagnosing mental health disorders. The questionnaires are primarily used in identifying the severity, too. Unlike physical illnesses, mental health diseases do not usually have objective biological markers, which makes the diagnosis much more dependent on the interpretation and judgment of the doctor.

Primary challenges in diagnosing hard-to-diagnose diseases:

Many mental health disorders are often considered complex/hard to diagnose disorders; this is basically because these kinds of disorders share common symptoms and they overlap with each other. Some examples of overlapping symptoms are, for eg, fatigue, sleep disturbance, distress, etc. Mental health disorders are expressed through behavior, speech, and expressions, but each of them is subjective as it varies from culture to culture.[10]

Limitations of human diagnosis:

Clinical diagnosis faces several systemic challenges. First, assessments are time-intensive and require sustained interaction. Second, there is a global shortage of trained professionals, particularly in resource-constrained settings. Third, financial barriers and social stigma reduce accessibility. These limitations highlight the need for scalable systems that can assist clinicians in early-stage screening.

How can AI be incorporated in the diagnosis of mental health?:

Artificial intelligence offers amazing tools for analyzing mental health problems. These tools include the use of automated, continuous, and consistent processing of large volumes of text provided by the patient.[11] Artificial intelligence for mental health and mental illnesses: an overview." Recent advances in machine learning and natural language processing have enabled AI models to analyze written and spoken language. AI models can analyze textual data such as clinical notes, questionnaires, and patient-reported outcomes. AI-powered chatbots can offer non-stop support to a larger demographic area at a reduced cost.[12]

AI-integrated mental health diagnosis:

Existing works in diagnosing mental health with the help of AI:

Currently, clinical notes, self-report questionnaires, online forums, and chatbot conversations have been widely used to identify mental health signals using machine learning and NLP models. It can be used in screening for depression, detecting suicide risks, and monitoring anxiety. Such systems usually formulate mental health screening as a task of supervised classification with labeled datasets.[13]

Types of input data for AI screening:

Clinical notes: contain observations recorded by a doctor during the interview. NLP models can extract useful information (symptom indicators, diagnostic cues) and use it to diagnose the disorder.[8]

Self-report questionnaires: This type of input data includes standardized symptom measures such as PHQ-9, and these are widely used for the process of training an AI.

Online text and social media: online posts from blogs and social media platforms capture a variety of expressions of distress and emotional states. These data are large-scale and reflect natural language use but are often context-dependent. Linguistic ambiguity and population bias are major challenges with this input type currently.

Model Architectures Used in Prior Studies:

Traditional machine learning models:

The most common traditional methods to analyze mental health data involve manually creating natural language processing based on simple pre-defined linguistic characteristics, such as term frequency (TF-IDF), word sentiment, or syntax (parts of speech. Traditional statistical methods that have been used in machine learning models include logistic regression, SVMs (Support Vector Machines), and Random Forests. However, their reliance on predefined features limits their ability to capture contextual meaning. In cases of textual ambiguity—where symptoms are expressed indirectly or overlap across multiple conditions—these models fail to distinguish between nuanced expressions. For example, phrases like "I feel off" or "something isn't right" lack explicit markers and are often misclassified or ignored entirely. As a result, traditional models are effective only when symptoms are explicitly stated, but they struggle significantly with subjective and context-dependent language.[14] This workflow is summarized in Figure 1.

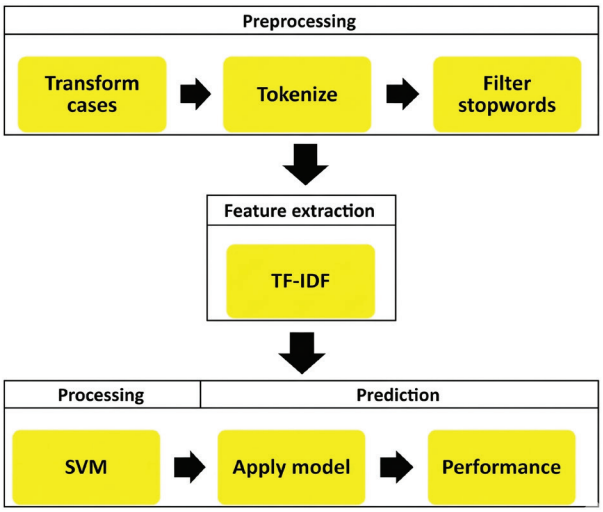


Figure 1: Figure 1 shows a traditional machine learning pipeline for mental health screening. It moves from preprocessing and TF-IDF feature extraction to an SVM model and performance evaluation.

Deep learning models:

Deep learning models, including convolutional neural networks, learn representations directly from raw text. They can capture local patterns such as key symptom phrases, improving performance over traditional methods. However, these models remain sensitive to variability in data and often fail to generalize when linguistic patterns deviate from training distributions, reducing their effectiveness in real-world scenarios.[15] A simplified version of this process is presented in Figure 2.

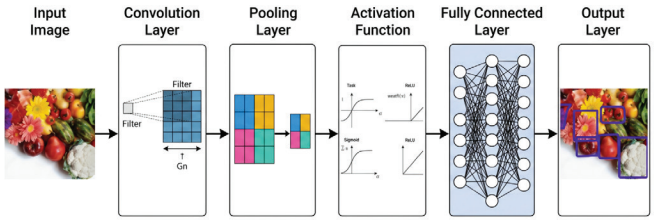


Figure 2: Figure 2 shows a deep learning classification architecture. It demonstrates how an input is processed through convolution, pooling, activation, and fully connected layers before producing an output prediction.

Transformer-Based Models (BERT, GPT, etc., for example):

Transformer-based models, such as BERT, represent the current state of the art in NLP. These models capture contextual relationships across entire text sequences using attention mechanisms.

Despite their advantages, transformers are not inherently equipped to handle ambiguity. Their performance depends heavily on training data quality. When inputs contain conflicting or incomplete symptom descriptions, model predictions can become unreliable.

For example, in cases where symptoms overlap between disorders (e.g., anxiety and PTSD), transformers may assign a single dominant label rather than reflecting uncertainty or

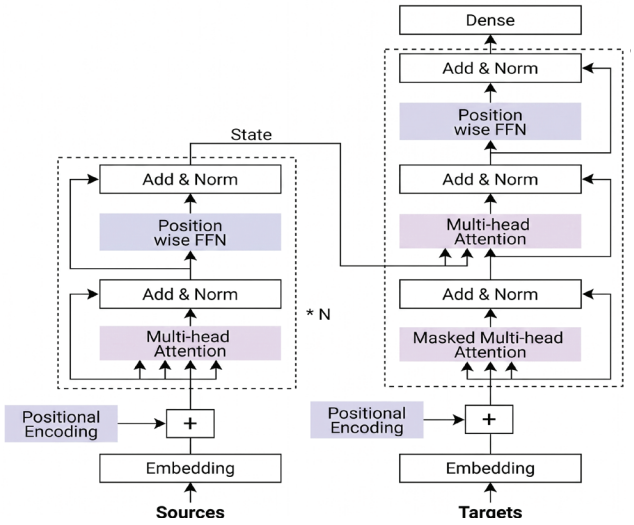


Figure 3: Figure 3 shows the basic transformer architecture, including embeddings, positional encoding, multi-head attention, feed-forward layers, normalization, and dense output. This helps explain why transformer models can capture context but may still struggle when patient language is incomplete or ambiguous.

■ Evaluation metrics used in previous studies

The majority of the studies test AI systems in terms of conventional performance measures like accuracy, precision, recall, F1 score, and PR-AUC. Although these measures are used to measure technical performance, they are not useful in measuring usability or trust. The recent studies focus on human-centered assessment with measures on a Likert scale for clarity, usefulness, appropriateness, and reliability. It is becoming evident in the literature that high accuracy is not the end goal for clinical deployment. The systems should also be able to express uncertainty and match the expectations of the user, especially in cases of ambiguous disorders.

■ Future worksss

Improved model architectures for complex disorders:

Future studies need to develop models of different architectures to take into consideration the ambiguities in language used to describe the symptoms and the overlap in symptoms of disorders, and not treat the process of screening for mental health as a typical classification problem. The majority of today's systems are optimized for disorders that have clear and

DOI: 10.36838/IJHSR820.9

standardized lists of symptoms and therefore fail miserably when attempting to classify disorders based on ambiguous, indirect, or heterogeneous language. Although transformer-based architectures have shown superior understanding capabilities, they also continue to have difficulties in interpreting culturally based symptom narratives. Therefore, future architectures that will be able to screen for complex disorders such as PTSD and depression subtypes must integrate representation learning with uncertainty modeling and structured symptom reasoning.

Multiple dataset integration:

While this review does focus mainly on text-based screening, future systems could benefit a lot from different data integration. When we combine patient language with physiological or behavioral signals, such as sleep patterns and heart rate, the variability may provide additional context that can help the AI system by reducing diagnostic ambiguity. Multimodal approaches have been suggested as a strong way to strengthen AI screening systems by compensating for the limitations of subjective language alone. Such integration should be carefully designed to respect the patient's privacy and ensure that non-textual signals are meaningful.

Human + AI collaboration in Mental health disorder diagnosis:

Future systems should adopt a human–AI collaborative model rather than fully automated diagnosis. In this framework, AI provides screening, risk scales, and symptom summaries, while clinicians will be responsible for final diagnostic decisions. Studies evaluating mental health chatbots show that user trust and perceived usefulness depend strongly on clarity, appropriateness, and safety rather than technical accuracy alone. A collaborative framework ensures that AI augments clinician judgment instead of replacing it and allows ambiguous cases to be escalated to professional evaluation.[17] The collaborative workflow proposed for this issue is shown in Figure 4.

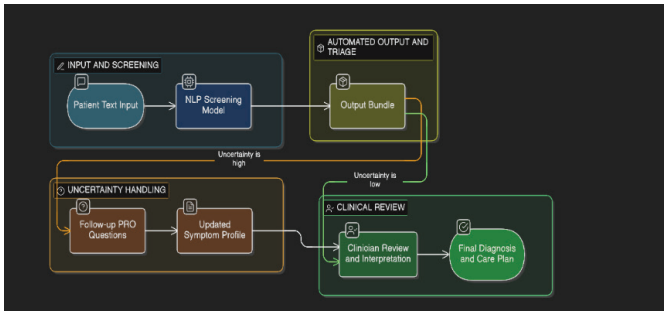


Figure 4: Figure 4 shows the proposed human-AI diagnostic workflow. Patient text is screened by an NLP model, uncertainty triggers follow-up PRO questions, and high-risk or unclear cases are routed to clinician review.

■ Proposed AI System Architecture

To address the limitations identified, this research proposes a conceptual AI architecture specifically engineered for the diagnosis of textually ambiguous mental health disorders.

The system consists of the following integrated components:

- Refined text input processing
- Contextual embedding utilizing transformer-based models
- Specialized ambiguity detection module
- Multi-label prediction layer
- Uncertainty estimation and probabilistic reasoning
- Human–AI interaction interface

Core Innovation:

The primary innovation of this architecture lies in its explicit modeling of linguistic ambiguity. Instead of dismissing vague patient inputs as noise, the system actively detects and quantifies uncertainty.

Multi-Label Prediction:

The framework generates multiple potential diagnoses rather than just a singular label. It mirrors real-world clinical reasoning in cases where symptoms significantly overlap across different mental health conditions.

Uncertainty Estimation:

Detailed confidence scores are produced for each diagnostic prediction, which allows the system to transparently communicate the degree of reliability associated with its outputs.

Human–AI Collaboration:

Diagnostic insights are presented to clinicians alongside confidence levels and ambiguity indicators, and they ensure that final interpretive decisions remain under professional human supervision. The final architecture is summarized visually in Figure 5.

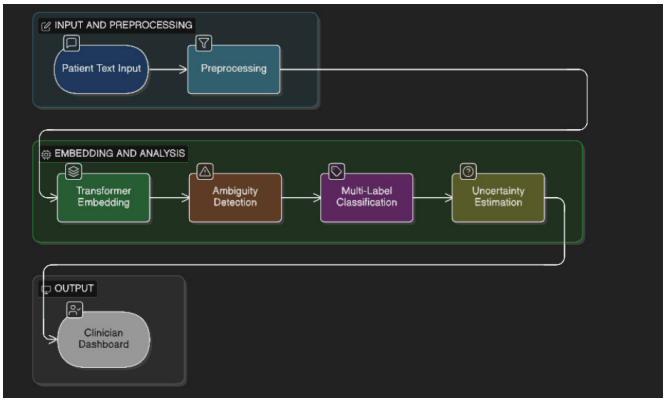


Figure 5: Figure 5 shows the proposed uncertainty-aware AI system architecture. It combines text preprocessing, transformer embeddings, ambiguity detection, multi-label classification, uncertainty estimation, and clinician-facing output.

■ Conclusion

This paper addresses the question of how AI systems can be strengthened to aid the screening of textually ambiguous mental health disorders using varying patient-reported outcomes and argues that this progress requires moving beyond accuracy-focused classification toward uncertainty-aware, practical, and much more human-centered design(one that can excel in understanding contexts). By using prior work across input data types, model architectures, and evaluation frameworks, this review highlights the gap between successful applications in

common disorders and the ongoing challenges associated with complex conditions such as PTSD and depression subtypes. The proposed focus on combining traditional performance metrics with a Likert-scale human evaluation aligns with the emerging healthcare AI frameworks that prioritize safety, transparency, and clinical collaboration. Ultimately, responsible deployment of AI in mental health care depends not on replacing clinicians, but on diagnostic reasoning and expanding access to early screening in a manner that acknowledges linguistic ambiguity and patient subjectivity. This research is limited by its methodological reliance on secondary literature synthesis and illustrative experimentation. Future work in this domain should prioritize validation utilizing diverse real-world datasets and rigorous clinical evaluation to ensure the practical efficiency of diagnostic frameworks.

■ Acknowledgments

I want to acknowledge my mentor, Dr Siddharth Krishnan, and my TA, Samuel Lefcourt.

■ References

1. Mental disorders. https://www.who.int/news-room/fact-sheets/detail/mental-disorders (accessed 2026-02-20).
2. Mental Health ATLAS 2020. https://www.who.int/publications/i/item/9789240036703 (accessed 2026-02-20).
3. D'Alfonso, S. AI in Mental Health. Curr. Opin. Psychol. 2020, 36, 112–117. https://doi.org/10.1016/j.copsyc.2020.04.005.
4. Is AI the Future of Mental Healthcare? | Topoi | Springer Nature Link. https://link.springer.com/article/10.1007/s11245-023-09932-3 (accessed 2026-02-20).
5. Artificial Intelligence for Chatbots in Mental Health: Opportunities and Challenges | Springer Nature Link. https://link.springer.com/chapter/10.1007/978-3-030-67303-1_10 (accessed 2026-02-20).
6. Schnittker, J. The Diagnostic System: Why the Classification of Psychiatric Disorders Is Necessary, Difficult, and Never Settled; Columbia University Press, 2017.
7. The Rise of Artificial Intelligence in Healthcare Applications. In Artificial Intelligence in Healthcare, Academic Press, 2020; pp 25–60. https://doi.org/10.1016/B978-0-12-818438-7.00002-2.
8. Natural language processing applied to mental illness detection: a narrative review | npj Digital Medicine. https://www.nature.com/articles/s41746-022-00589-7 (accessed 2026-02-21).
9. Regier, D. A.; Kuhl, E. A.; Kupfer, D. J. The DSM-5: Classification and Criteria Changes. World Psychiatry 2013, 12 (2), 92–98. https://doi.org/10.1002/wps.20050.
10. Liu, J.; Wang, H.; Li, L. Rethinking the Studies of Diagnostic Biomarkers for Mental Disorders. Meta-Radiol. 2025, 3 (1), 100135. https://doi.org/10.1016/j.metrad.2025.100135.
11. Artificial Intelligence for Mental Health and Mental Illnesses: an Overview | Current Psychiatry Reports | Springer Nature Link. https://link.springer.com/article/10.1007/S11920-019-1094-0 (accessed 2026-02-21).
12. Arji, G.; Erfannia, L.; Alirezaei, S.; Hemmat, M. A Systematic Literature Review and Analysis of Deep Learning Algorithms in Mental Disorders. Inform. Med. Unlocked 2023, 40, 101284. https://doi.org/10.1016/j.imu.2023.101284.
13. Analysis of depression in social media texts through the Patient Health Questionnaire-9 and natural language processing - Nam Hyeok Kim, Ji Min Kim, Da Mi Park, Su Ryeon Ji, Jong Woo Kim, 2022. https://journals.sagepub.com/doi/full/10.1177/20552076221114204 (accessed 2026-02-21).
14. A novel text mining approach based on TF-IDF and Support Vector Machine for news classification | IEEE Conference Publication | IEEE Xplore. https://ieeexplore.ieee.org/abstract/document/7569223 (accessed 2026-02-21).
15. Alzubaidi, L.; Zhang, J.; Humaidi, A. J.; Al-Dujaili, A.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M. A.; Al-Amidie, M.; Farhan, L. Review of Deep Learning: Concepts, CNN Architectures, Challenges, Applications, Future Directions. J. Big Data 2021, 8 (1), 53. https://doi.org/10.1186/s40537-021-00444-8.
16. End-to-End Transformer-Based Models in Textual-Based NLP. https://www.mdpi.com/2673-2688/4/1/4 (accessed 2026-02-21).
17. JMIR mHealth and uHealth - Validity of Chatbot Use for Mental Health Assessment: Experimental Study. https://mhealth.jmir.org/2022/10/e28082 (accessed 2026-02-21).

■ Author

Dhairya Jain is a Grade 11 student at Jayshree Periwal Global School, Jaipur. He plans to pursue Computer Science, with interests in how artificial intelligence and technology can be used to solve real-world problems. He actively engages in research, Model United Nations, and drama, combining technical curiosity with leadership and communication.

DOI: 10.36838/IJHSR820.9

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

How Neurotransmitters Shape Depression

Dhea Katkuri

Alliance Academy for Innovation, 1100 Lanier 400 Pkwy, Cumming, GA 30040, USA; dhearkatkuri@gmail.com

ABSTRACT: Depression has traditionally been explained as a chemical imbalance in the brain. However, growing evidence suggests that it involves interactions among neurotransmitters, brain networks, neuroplasticity, and environmental factors. Even though the chemical imbalance model has contributed to expanding knowledge concerning depression, this theory does not offer an explanation that can help account for the various reactions exhibited by different patients. More research is therefore needed to understand how neurotransmitters affect stress. This study analyzed the roles of serotonin, dopamine, norepinephrine, and glutamate in depression through a literature review of scientific journal articles. This review revealed that depression is understood as a disorder involving disruptions across interconnected brain systems rather than a deficiency of a single neurotransmitter. It has been determined that these neurotransmitters control the process of emotion, motivation, attention, the stress response mechanism, and neuroplasticity. Chronic stress triggers dysfunction of the hypothalamus-pituitary-adrenal axis, elevated concentrations of cortisol, reduced neurogenesis, and low BDNF levels. These findings suggest that depression develops through interactions among neurochemical, structural, and environmental influences. Understanding depression from an integrated systems perspective may support more personalized treatment approaches that combine medication, psychotherapy, and lifestyle interventions.

KEYWORDS: Cellular and Molecular Biology, Neurobiology, Depression, Neurotransmitters, Brain-derived Neurotrophic Factor (BDNF).

Introduction

For many decades, scientists have explored various areas of the brain in an effort to discover how different chemicals affect emotions and behavior. Serotonin is one of the most widely studied neurotransmitters that helps to regulate emotional responses. When serotonin signaling is impaired, people are more prone to experience depression and anxiety.[1,2] Similarly, dopamine is an important neurotransmitter, responsible for regulating motivation, reward, and feelings of pleasure. Low levels of dopamine have been associated with anhedonia, which is the inability to feel pleasure, a symptom of depression.[3,4] Another neurotransmitter, norepinephrine, is responsible for maintaining alertness and focus in a person and has been linked to fatigue and lack of concentration when not functioning properly.[5] Moreover, glutamate is an essential neurotransmitter involved in the process of learning and memory formation in the brain. In recent studies involving ketamine administration, glutamate was also found to have mood-stabilizing and plasticity properties.[6,7] All these neurotransmitters work together and influence emotion, motivation, attention, learning, and behavior at the same time.[3,8]

With a better understanding of how the brain works, scientists have begun to question the idea that depression can be caused by a specific imbalance within the brain. Neurotransmitters, for example, do not work independently; rather, they affect each other and other systems in the body. Moncrieff *et al.* discovered that the data on the cause of depression is inconclusive since there was no scientific proof linking depression to the lack of serotonin levels in the brain.[2] Other research links problems in neurotransmitters such as dopamine and norepinephrine with symptoms including low energy, low motivation, and fatigue.[3,9] Nevertheless, even though these studies clarify some characteristics of depression, they cannot account for the variety of symptoms experienced by patients. Therefore, researchers describe depression as a condition affected by neurotransmitters, stress hormones, neuroplasticity, and communication among different brain regions.[3,8]

More recently, evidence for this broader perspective has been strengthened by large-scale brain networks, including the Default Mode Network (DMN). Studies have found that individuals with depression often show altered connectivity involved in self-reflection, emotional processing, and cognitive control.[8,10] Increased activity within these networks has been linked to rumination, the repetitive negative thinking that commonly accompanies depression.[8] These findings suggest that mental illness is influenced not only by the presence of specific neurotransmitters but also by how different brain regions communicate with one another. Neuroplasticity is an important part of this process because the brain's connections can strengthen, weaken, and reorganize in response to life experiences, environmental influences, and treatment.[8,11]

Depression is also caused by chronic stress through the impact of stress on the hypothalamic pituitary adrenal (HPA) axis. Pariante and Lightman revealed that long-term stress can cause an impairment of the functioning of the HPA axis, increase cortisol secretion, and alter brain functioning.[9] Sapolsky also noted that chronic exposure to glucocorticoids could be implicated in the shrinkage of the hippocampus and reduction in the number of neurons that could lead to the development of depression.[12] Another factor related to depression is Brain-Derived Neurotrophic factor. This protein, which stimulates neurons' growth and viability, has been implicated as a signif-

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.10

icant player in mood disorders. Depression could result from decreased BDNF activity, whereas successful therapy leads to increased BDNF activity.[11]

This paper views depression not as a disorder caused by a single chemical imbalance, but as a condition involving multiple interacting brain systems. By examining neurotransmitter signaling, HPA-axis regulation, neuroplasticity, and brain network function, this paper explores the neurobiological mechanisms that contribute to depression. Looking at depression from this broader perspective helps explain why different people respond to different treatments. Medication, therapy, lifestyle changes, and stress management techniques may each target various parts of the same interconnected system. Research that demonstrates the effectiveness of antidepressants, psychotherapy, and emerging glutamatergic treatments further supports the need for a comprehensive strategy to treatment.[6,7] Understanding these relationships may help researchers and clinicians develop more personalized and effective approaches to care.

Methods

This paper reviews the roles of serotonin, dopamine, norepinephrine, and glutamate in depression. Sources for this review were identified through searches of PubMed, Google Scholar, and library databases. Search terms that were used were "depression serotonin," "depression HPA axis," "depression BDNF," "depression glutamate," "depression reward system," "major depressive disorder dysregulation of networks," and "the involvement of dopamine in depression."

This review is a narrative review rather than a systematic one because the goal was to synthesize the major findings across relevant studies, rather than to obtain every study that could be found using a meta-analysis. Sources were eliminated if they were not peer-reviewed, were mostly opinionated, were not related to depression, or did not directly speak to neurotransmitters, neural circuits, biology of stress, or depression.

A total of 18 primary sources including peer-reviewed journal articles, clinical trials and meta-analyses were reviewed based on their relevance, scientific significance, recency (2000-2026), and contribution to understanding the "chemical imbalance model" including perspectives that both challenged and expanded upon traditional neurotransmitter-based explanations, such as the review by Moncrieff *et al.*[2] In addition to neurotransmitter alterations, this review examined the roles of neuroplasticity, the HPA axis, and BDNF, all of which help mediate the effects of stress and the environment on the brain.[9,11]

The DSM-5-TR definition of major depressive disorder was used to maintain diagnostic consistency in the reviewed literature.[13] Most importantly, the study organized the results according to various functions served by these systems, such as affective regulation, motivation, and stress response, while noting the interactions among the systems, instead of isolating each from the other.[4,5]

Results

Serotonin and Emotional Regulation Serotonin (5-HT) regulates the sensitivity, timing, and strength of neural responses to emotional stimuli.[2] It is not a fast-acting, localized generator of positive mood, but is more likely to modulate the strength, timing, and ordering of emotive responses. In this way, it may be thought of as a kind of global filter across the brain. 5-HT is produced from the amino acid tryptophan in the raphe nuclei of the brainstem, which project serotonergically to virtually every forebrain structure, including the prefrontal cortex, amygdala, hippocampus, and hypothalamus. The scope of its projection system allows wide-ranging influences on behavior and physiology, ranging from mood to impulse control to sleep.[14] Following receptor activation, serotonin is taken up by the 5-HT transporter into the presynaptic neuron, and the availability of extracellular 5-HT is controlled by reuptake efficiency.[2,14]

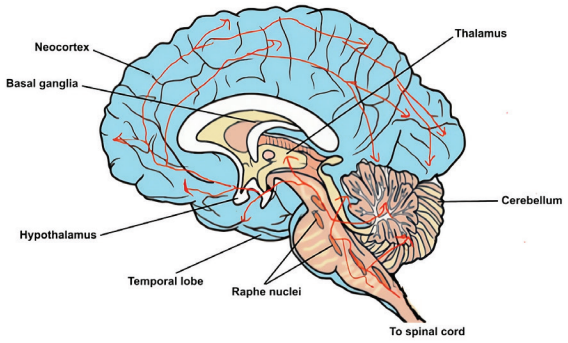


Figure 1: Simplified Serotonergic Pathways. Serotonin neurons originate primarily in the raphe nuclei and project throughout the brain, including the neocortex, thalamus, hypothalamus, and basal ganglia. The distribution of these pathways enables serotonin to influence mood, cognition, sleep, and other physiological functions. Redrawn from ref.[18]

As shown in Figure 1, serotonin pathways, which arise from the raphe nuclei, travel to different brain regions. Since serotonin pathways extend throughout the brain, it affects emotions, cognition, sleep, appetite, and physiological homeostasis.[14] Different types of receptors are activated by serotonin, with the 5-HT1A receptor, for instance, tending to reduce activity whilst calming the emotions, and the 5-HT2A receptor increases activity in the cortex and has a central role in perception and cognition.[14] This receptor heterogeneity enables serotonin to precisely modulate neural signaling in the brain without artificially boosting or reducing activity. So, reduction of their levels or receptor effectiveness will have reduced reliability in controlling the emotionally significant stimuli.[5,14] It can be deduced that a reduction in serotonergic signaling might impair the brain's ability to regulate the emotional salience of stimuli, leading to a relatively higher attentional and memory-based salience of adverse events, consistent with the body of evidence for negative bias and emotional cognitive processing in depression.[5,15]

For instance, a loosening of regulatory control may give rise to repetitive thoughts and worry, whereby small negative events are woven into a narrative that supports bigger (moral) failures. At the same time, a diminished restraint of the amygdala will intensify emotionality.[4,8] Changes will also be

registered in eating and sleeping behaviors due to the implication of serotonin in circadian and hypothalamic processes.[15] More serotonin in the brain does not necessarily equate with happiness; it may soon enhance the brain's capacity to more gradually maintain emotion and related behavior over time, and so return the affect regulation to a state such that the brain more readily manages that affectively salty moment without stressors generating stress.[5,8]

With more reality-based amygdala activation and more powerful and effective PFC influences, a realm of optimal cognition and comportment, the neural states more gracefully hip up (prescribe a more context-appropriate state).[8,14] An overabundance of serotonergic activity has the net effect of being effectively blunting or somnogenic.[14] Its mission would seem to be to elevate and maintain one's default homeostatic state when one is not actually performing, as much as possible, to optimize positive functioning.

Similar to other neurotransmitter systems (dopamine, noradrenaline, and glutamate), Serotonin interacts extensively with them. Such interaction has profound influential effects on motivation, arousal, and plasticity and suggests a possible role of serotonergic activity in the middle of a control station responsible for the regulation of cognitive and affective processes.[4,8] As a result, the extent to which serotonergic transmission is altered will influence perception, make sense of the environment, and respond to it. This further points to the pivotal role of serotonin and the neurobiology of depression.[3,4,15]

Dopamine and Reward Dysfunction:

Dopamine itself is a monoamine and is part of the reward system in the brain, yet its main interpretation of its action is in regard to the regulation of motivation. It influences motivation as to how motivated an individual will be to attain a certain goal and how much effort they will put in (motivation).[4,5] Dopamine is synthesized in the brain from a precursor amino acid, tyrosine, in cells by L-DOPA and is concentrated in neurons in the ventral tegmental area (VTA) and the substantia nigra.[14]

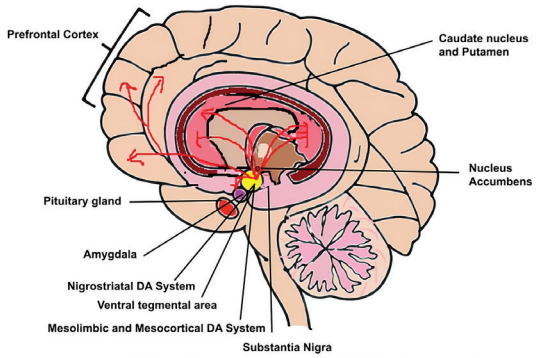


Figure 2: Simplified Dopaminergic Pathways. Dopaminergic neurons originate primarily in the ventral tegmental area and substantia nigra and project throughout the brain, including the nucleus accumbens, prefrontal cortex, amygdala, caudate nucleus, and putamen. These pathways enable dopamine to regulate reward processing, motivation, emotion, cognition, and motor function. Redrawn from ref.[16]

As shown in Figure 2, the dopaminergic pathways arising from these areas, indicated by the arrows, project to the prefrontal cortex and the nucleus accumbens, and this is where the pathways for reward and its behaviorally related stimuli originate.[4,5,14] Dopaminergic fibers extending from these regions to the dorsal and ventral striatum constitute the mesolimbic and mesocortical pathways.[4,8,14] These pathways are responsible for the reward-related responses, decision-making, and behavior.[4,5]

Dopamine receptors within the synapses exist in two forms, the D1-like and D2-like receptors.[14] These are mostly excitatory, enhancing current cell signaling, or inhibitory, and modulate feedback within the reward processing pathways. This enables dopamine to control the relationship between expectation and reward and outcome-reinforcing behaviors. In fact, dopamine is not simply a pleasure transmitter; it is implicated in reward prediction error signaling, sending a message when outcomes are either better or worse than anticipated, thereby helping to update behavior. A system like this is necessary for learning, habit formation, and the ongoing pursuit of useful activities.[4,5,14]

In particular, a diminished dopamine mesolimbic transmission adversely affects the brain's capability to grant value to potential rewards. Reducing mesolimbic dopamine transmission may contribute to anhedonia, or a diminished ability to anticipate and experience reward.[4,5]

In this state, the brain's ability to generate strong predictive reward signals is reduced. For example, a student understands that studying is a way to success, but he or she has no motivation coming from an anticipated reward to start studying. In the end, depression levels increase as the rewarding experiences become less and less frequent, leading to withdrawal and inactivity, which further worsens the state of depression.[4,5]

In contrast, increases in dopamine can improve motivation, attention, and engagement in goal-related activities by making rewards feel more significant or worthwhile.[4,5] This increases the likelihood that a person will start and continue a task because the connection between effort and positive outcome feels stronger. However, excessively high dopamine activity can also disrupt the brain's reward system and contribute to impulsive behavior, risk-taking, or addiction. This suggests that dopamine's role is not simply to increase pleasure, but to regulate motivation and reinforcement.[4,14]

Dopamine also works closely with other neurotransmitters like serotonin, norepinephrine, and glutamate.[4,8] Together, these systems regulate mood, manage alertness, and support cognitive flexibility. Because of this interaction, dopamine becomes part of the larger neural network involved in motivation and behavior. When these systems become dysregulated, depression can affect not only mood but also a person's ability to take action, pursue goals, and find meaning in daily life. This helps explain why dopamine is considered a key factor in the neurobiology of depression.

Norepinephrine and Stress Response:

Norepinephrine (noradrenaline/NE) is a monoamine neurotransmitter and hormone that governs waking stress, attention, and levels of arousal. For it to be able to govern different areas of the brain, NE must map out to different areas.

DOI: 10.36838/IJHSR820.10

The neurotransmitter is synthesized from dopamine by the dopamine hydroxylase enzyme and is mainly produced by the locus coeruleus cells.

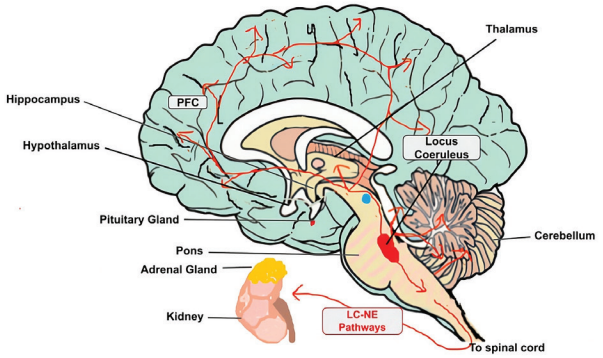


Figure 3: Simplified Locus Coeruleus Pathways. Noradrenergic neurons originating in the locus coeruleus project extensively throughout the brain, including the prefrontal cortex, hippocampus, hypothalamus, and thalamus, and interact with the HPA axis involved in the stress response. The widespread distribution of these pathways enables norepinephrine to regulate arousal, attention, cognition, and emotional processing. Redrawn from ref.[17]

As shown in Figure 3, the locus coeruleus cells have a wide and dense projection of the norepinephrine system throughout the brain. These give off axons to target different areas of the brain, including the prefrontal cortex, amygdala, hippocampus, and thalamus. This wide projection network is thought to enable the coordination of waking/cognition/emotion.

Norepinephrine acts at a synaptic level by way of adrenergic receptors, which are separated into alpha and beta types that influence neurons in various ways.[14] Alpha-2 receptors tend to be inhibitory autoreceptors that stop norepinephrine release and restore normal function, whereas Alpha and Beta receptor types are typically excitatory. Through this receptor system, norepinephrine influences the brain's preparedness to react to stimuli, thereby regulating the degree of alertness and attention per the needs of the environment.

Unless the brain has enough norepinephrine, a person cannot stay awake and focused. Fatigue, slow thinking, and inability to concentrate are all symptoms that are caused by underactive brain cells that help us pay attention. This is sometimes called "mental fog", when even the simplest things seem exceedingly difficult. For instance, a student could find it extremely hard to do homework, not because he or she does not know the material, but because the brain cannot stay alert to the point that it could fully understand the material. Additionally, less norepinephrine means that the brain becomes less good at stress response, which results in less motivation and a lower level of reaction to external changes.[5,14]

Norepinephrine increases arousal, alertness, and cognition when there is a need for quick reaction or sustained concentration. Moderate increases of Norepinephrine enhance concentration and goal-directed behavior in the prefrontal cortex. Excessive amounts of norepinephrine are indicative of hyperarousal due to chronic stress, which appears as anxiety, irritability, and "a superdominance" of threat-detecting systems. As the system becomes oversensitive to potential threats, environmental stimuli elicit exaggerated threat-detection signals, resulting in sleep dysfunction, "bad" decision-making, and emotional lability. These six account for the comorbidity of depression and anxiety.[5,9,14]

Comprehensive interactions exist between norepinephrine and the serotonin, dopamine, and glutamate systems, largely related to control of the hypothalamic, pituitary, and adrenal axis, which orchestrates the body's response to stress. Chronic activation causes increased cortisol levels, which alter norepinephrine activity, resulting in neuronal dysregulation. Norepinephrine acts as the major intermediary between physiological stress and cognition/emotion, raising the possibility that dysregulation causes not only a mood disturbance but a dampening of attention, pervasive fatigue, and physiological exhaustion.[8,9]

Glutamate and Emerging Research:

Glutamate acts as the main excitatory neurotransmitter within the brain, facilitating rapid synaptic transmission and synaptic plasticity; in contrast to monoamine neurotransmitters, which were observed to subtly change the overall 'mood' of the brain, glutamate was observed to directly facilitate long-term potentiation or depression, and the formation or retraction of synapses. This is how the brain learns and forms memories. At the molecular level, glutamate has an effect on five receptor types. Most relevant in the CNS are N- methyl-D-aspartate (NMDA) and α-amino-3-hydroxy-5-methyl-4-isoxazolepropionicacid (AMPA) receptors. Such stimulation of glutamate receptors affects the opening of calcium channels and elicits diverse signaling pathways leading to changes in gene expression, synaptic growth, and long-term potentiation. Long-term potentiation is hypothesized to be the cellular explanation for learning through repeated strengthening of synaptic activity.[14]

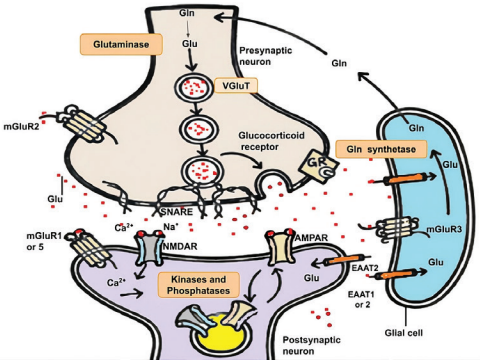


Figure 4: Simplified Glutamate-Glutamine Neurotransmitter Cycle. Glutamate released from the presynaptic neuron activates NMDA and AMPA receptors on the postsynaptic neuron, contributing to excitatory neurotransmission and synaptic plasticity. Excess glutamate is removed by astrocytic transporters and recycled through the glutamate-glutamine cycle, helping maintain neurotransmitter balance. Redrawn from ref.[14]

As shown in Figure 4, when glutamate signaling pathways become deregulated, plasticity decreases, compromising the connectivity between neural circuits, causing depression. Chronic stress also serves as one of the external factors that lead to depression, provoking an excessive release of glutamate in some brain regions, mainly in the prefrontal cortex. This excessive stimulation of calcium influx causes excitotoxicity, leading to neuronal damage, specifically characterized by den-

dritic atrophy and a decrease in synaptic density. This leads to slower communication between areas involved in working memory and emotional regulation. The hippocampus, on the other hand, shows opposite responses, such as a decrease in glutamatergic activity and formation of synapses.[7,11,14]

Depression causes physical damage to the brain tissue, which weakens the functional output. This is true within the hippocampus, where the brain's memory mediator and "reality sense" is to a larger extent than in several other parts of the brain, which explains the memory lapses often associated with depression; the brain is simply not capable of recording the experience, leaving the individual unable to remember anything new, the way in which it loses effectiveness and lacks alertness. With glutamate affecting what can only be called a blow to the brain's ability to affect itself, depression is a state of mental hibernation, immobile and depressed negativity.

Though central to bringing about peak functioning through reinstatements or replacements of glutamate, the transmission process takes hours. One example is the use of drugs like ketamine, an NMDA (N-methyl-D-aspartate) antagonist. By blocking NMDA receptor activity and triggering a cascade that enhances AMPAR transmission and neurotrophic factor expression, ketamine may promote synaptic plasticity, a mechanism that was proposed by Berman et al and Krystal et al to contribute to its rapid antidepressant effects.[7,11] This fast and structural resolution of depression is not because its victims have had their own feelings directly ignored, but because their own cognition has been put into a functional state. It is this way that glutamate appears to be more than a facilitator over the poles of emotion: it is fundamental to the very mental mechanics through which individuals live in emotion. When glutamate regulation fails, depression becomes not just an emotion, but a structural breakdown of the brain.

Neuroplasticity and Brain Structure:

Though each neurotransmitter does necessarily have a different function, depression itself is a neuroplasticity disorder, a deficient adaptation in plasticity; the brain's ability to restructure itself.[11] Brain-derived neurotrophic factor (BDNF) is a key biological mediator of neuroplasticity and makes neuronal cell survival, growth, and synaptic communication easier.

Normally, through the action of BDNF, new neurons in the hippocampus are created in a process called neurogenesis. BDNF also aids in making synaptic connections stronger, which is a feature of learning and is also influenced by the environment. However, persistent stress and neurotransmitter imbalance for a long duration cause a drop in BDNF levels, and that diminishes the brain's ability to change and evolve.[9,11]

This reduction is partly driven by sustained activation of the hypothalamic-pituitary-adrenal axis, which leads to repeated cortisol release. In addition to being the body's primary stress hormone, elevated cortisol can suppress BDNF production and contribute to structural changes in the brain. The hippocampus undergoes a slight reduction in size over time, while the connectivity and functioning of the prefrontal cortex are also impaired.[8,10,12] These neuroanatomical changes do not merely correlate with the aging process but have functional consequences.[3] An under-sized hippocampus compromises context processing, which results in significant difficulty in differentiating between now and then, which threatens to continue negative effects by impairing the brain's ability to update its internal model of reality.[8,12]

Similarly, when the prefrontal cortex becomes dysfunctional and cannot top-down control the brain's emotion reactivity hub, the amygdala, emotions will probably be less self-regulated and heightened. In fact, this would cause states of heightened emotion, in particular under duress or threat. Yet, affliction by 'depression' is not a cause for permanent change. Therapeutic intervention through treatment that brings neurotransmitter and diathesis parameters to 'normal' levels reduces the stress elicited and increases BDNF production and neurogenesis. Both neuroscience and psychological techniques have shown that these brain 'connectivities' can be enhanced, showing that the 'depressed' brain is capable of being transformed.[10,11]

Integrated Systems Perspective: Network Dysregulation Rather Than Chemical Deficiency:

Modern neuroscience is moving away from the "one chemical" explanations for depression and is considering it a disorder of the whole brain, involving network dysfunction rather than one single neurotransmitter. Neurotransmitters do not operate in isolation, but are part of complex networks that govern cognition, emotion, and behavior. The problem is no longer one of "low" chemicals, but impaired messaging and interaction between brain regions. Another important set of pathways involves the interpretation of individual experiences.

Serotonin is involved with experience interpretation, dopamine with the perception of reward, norepinephrine with arousal and stress responses, and glutamate with neural plasticity.[3-5,8] These pathways are constantly in a feedback loop with each other to either support good mental health or create a disturbed mentality.

An example of the effects of stress can be seen where the spikes in the level of cortisol dysregulate the function of norepinephrine and prompt an oversecretion of glutamate. Not only does this inhibit synaptic connectivity, but it also lowers BDNF levels, which impairs neuroplasticity, leading to a lack of control of emotion by serotonin and a lack of motivation by dopamine.[9,11]

Given that depression is the result of the interaction of poor functioning of various systems, treatment must not only involve treating one neurotransmitter. For instance, traditional antidepressants such as selective serotonin reuptake inhibitors (SSRIs) and selective norepinephrine reuptake inhibitors may operate by gradually altering the signaling between serotonin and norepinephrine transmission. Certain drugs like ketamine and other recently discovered drugs in this class operate directly on Glutamate levels and establish a quick increase in nerve networks. Other SSRIs alone are inadequate, as the condition has an implication on conduct, thinking, and surroundings, too. Adjustments to the conduct and surroundings may prove particularly useful in reforming nerve cell function.

For example, apart from the increase in BDNF, regular physical activity also helps with neurogenesis and optimum

DOI: 10.36838/IJHSR820.10

sleep-wake rhythms, maintains daily circadian rhythms, and supports neurotransmitters' circadian synchronization. On the other hand, cognitive behavioral psychotherapy brings intervention for training brain circuitry to function differently by interrupting negative thinking and reinforcing the prefrontal regions involved in emotion regulation.

These treatments do more than simply increase neurotransmitter levels. Over time, they help restore communication between neurons, allowing the brain to function more effectively. As these connections recover, the brain becomes better at adapting to challenges, processing emotions, and responding to the world around it.

This model accounts for observed differential symptom dimensions in depression across people. One person may experience only fatigue and anhedonia, another anxiety and hyperarousal, or another rumination and rigidity of cognition. These differential symptoms may prove to be certain dysfunctional nodes in a common network of neurotransmitters and neural systems. Overall, depression as a disorder of network disruption can elucidate the how and the why about the etiology and treatment of depression.

■ Discussion

Depression is not a single condition influenced by neurotransmitters or certain brain circuits, but it is a phenomenon resulting from neurophysiology and environment as well.[9,11,15] To integrate information from various scientific fields such as neurobiology and medicine, this research attempts to study the role of serotonin, responsible for our emotions; dopamine, responsible for motivation and rewards; norepinephrine, influencing our ability to pay attention and cope with stress; and glutamate, involved in synaptic plasticity and adaptation.[4,7,15] It should be added that neurotransmitters do not operate in isolation from each other but influence brain activities through forming feedback loops so that any dysregulation triggers changes in the activity of other brain circuits via processes such as chronic stress, HPA axis activation, and deficiency of BDNF.[8,9,11] Thus, it can be said that the disorder of depression can be seen as one caused by dysfunction within the brain network, leading to emotional and mental changes that hinder actions, make people rigid, and unable to cope with stress.[4,8]

This understanding of depression may also help guide treatment and prevention strategies. Because depression affects multiple interacting brain networks, treatment should be tailored to the individual. Monotherapy may not be sufficient for all patients. Combining conventional drugs, such as SSRIs, SNRIs, and ketamine, with behavioral interventions such as cognitive behavioral therapy (CBT), regular exercise, and adequate sleep may provide better outcomes by supporting neuroplasticity, improving emotional regulation, and reducing depressive symptoms.

■ Conclusion

Depression was once mainly understood as a biochemical dysfunction. However, research now suggests that it cannot be explained by the dysfunction of a single neurotransmitter alone. When the communication between these brain circuits becomes disrupted, the brain struggles to manage stress, regulate emotions, and perform other normal cognitive functions. Serotonin, dopamine, norepinephrine, and glutamate do not work independently, but interact with one another through connected neural networks.

The adoption of such an understanding totally alters our conception of how we can treat the problem. The treatment for depression needs to go beyond the shortage of one chemical and incorporate several brain systems working simultaneously. Other interventions, such as clinical therapies, need to be supplemented by other approaches as well, such as healthy routines, socializing, and even a change in lifestyle. Depression is not just an emotional problem; it affects the brain's adaptation and response to change. Hence, the treatment should not focus on a single cause of depression but instead consider pharmacological, psychological, and behavioral factors that contribute to the disorder.

Future researchers should further investigate the mechanisms driving why individuals with depression often respond differently to the same treatment. Although many patients receive similar medications, some patients achieve full remission while others see no improvement. Examining genetic factors, environmental influences, and variations in brain network activity may help explain these disparities and support the development of more personalized approaches to treatment.

■ Acknowledgments

I would like to thank Professor Jobin Varkey for guiding me throughout this research project and helping me deepen my understanding of neurotransmitters and their impactful connection to depression. I also thank Professor Jen Tuban for reviewing and providing feedback that strengthened my work, and Coach Jo Kethar for keeping me motivated and providing me with these necessary resources. Finally, I owe my heartfelt thank you to my parents for giving me this opportunity to expand my knowledge on neuroscience as I continue to grow my interest and pursue medicine.

■ References

1. Caspi, A.; Sugden, K.; Moffitt, T. E.; Taylor, A.; Craig, I. W.; Harrington, H.; McClay, J.; Mill, J.; Martin, J.; Braithwaite, A.; Poulton, R. Influence of Life Stress on Depression: Moderation by a Polymorphism in the 5-HTT Gene. *Science* **2003**, *301* (5631), 386–389. DOI: 10.1126/science.1083968
2. Moncrieff, J.; Cooper, R. E.; Stockmann, T.; Amendola, S.; Hengartner, M. P.; Horowitz, M. A. The Serotonin Theory of Depression: A Systematic Umbrella Review of the Evidence. *Mol. Psychiatry* **2022**, *27*, 2397–2409.
3. Nestler, E. J.; Barrot, M.; DiLeone, R. J. Neurobiology of Depression. *Neuron* **2002**, *34* (1), 13–25. DOI: 10.1016/S0896-6273(02)00653-0
4. Russo, S. J.; Nestler, E. J. The Brain Reward Circuitry in Mood Disorders. *Nat. Rev. Neurosci.* **2013**, *14* (9), 609–625. DOI: 10.1038/nrn3381
5. Pizzagalli, D. A. Depression, Stress, and Anhedonia: Toward a Synthesis and Integrated Model. *Annu. Rev. Clin. Psychol.* **2014**, *10*, 393–423. DOI: 10.1146/annurev-clinpsy-050212-185606

6. Berman, R. M.; Cappiello, A.; Anand, A.; Oren, D. A.; Heninger, G. R.; Charney, D. S.; Krystal, J. H. Antidepressant Effects of Ketamine in Depressed Patients. *Biol. Psychiatry* **2000**, *47* (4), 351–354. DOI: 10.1016/S0006-3223(99)00230-9
7. Krystal, J. H.; Sanacora, G.; Duman, R. S. Rapid-Acting Glutamatergic Antidepressants: The Path to Ketamine and Beyond. *Biol. Psychiatry* **2013**, *73* (12), 1133–1141. DOI: 10.1016/j.biopsych.2013.03.026
8. Price, J. L.; Drevets, W. C. Neurocircuitry of Mood Disorders. *Neuropsychopharmacology* **2010**, *35* (1), 192–216. DOI: 10.1038/npp.2009.104
9. Pariante, C. M.; Lightman, S. L. The HPA Axis in Major Depression. *Trends Neurosci.* **2008**, *31* (9), 464–468. DOI: 10.1016/j.tins.2008.07.006
10. Schmaal, L.; Hibar, D. P.; Sämann, P. G.; Hall, G. B.; Baune, B. T.; Jahanshad, N.; Cheung, J. W.; van Erp, T. G. M.; Bos, D.; Ikram, M. A.; *et al.* Cortical Abnormalities in Adults and Adolescents with Major Depression. *Mol. Psychiatry* **2017**, *22* (6), 900–909. DOI: 10.1038/mp.2016.60
11. Duman, R. S.; Monteggia, L. M. A Neurotrophic Model for Stress-Related Mood Disorders. *Biol. Psychiatry* **2006**, *59* (12), 1116–1127.
12. Sapolsky, R. M. Glucocorticoids and Hippocampal Atrophy in Neuropsychiatric Disorders. *Arch. Gen. Psychiatry* **2000**, *57* (10), 925–935. DOI: 10.1001/archpsyc.57.10.925
13. American Psychiatric Association. *Diagnostic and statistical manual of mental disorders*: DSM-5-TR, 5th ed., text rev. American Psychiatric Publishing: Washington, DC, **2022**.
14. Stahl, S. M. *Stahl's Essential Psychopharmacology: Neuroscientific Basis and Practical Applications*, 4th ed.; Cambridge University Press: Cambridge, U.K., **2013**.
15. Belmaker, R. H.; Agam, G. Major Depressive Disorder. *N. Engl. J. Med.* **2008**, *358* (1), 55–68. DOI: 10.1056/NEJMra073096
16. Guy-Evans, O. Brain Reward System. *Simply Psychology*, **2026**. https://www.simplypsychology.org/brain-reward-system.html (accessed **2026-06-03**).
17. Driven Resilience. 3 Ways to Regulate Norepinephrine. **2018**. https://home.hellodriven.com/articles/neurohack-your-brain-for-resilience-3-ways-to-regulate-norepinephrine (accessed **2026-06-03**).
18. Shutterstock. *Serotonin System Diffuse Modulatory Systems Implicated*; Image ID 1621521616. https://www.shutterstock.com/ (accessed **2026-06-03**).

■ Author

Dhea Katkuri is a student at Alliance Academy for Innovation in Cumming, Georgia, interested in pediatric medicine, neuroscience, and mental health. She believes that the early years of life are most pivotal because they influence long-term neurological development. Dhea plans to pursue an undergraduate degree in the biology track.

DOI: 10.36838/IJHSR820.10

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Machine Learning for Heatstroke Prediction: Unlocking Weather and Climate Signals

Jingyao Zhu[1,2]

1) Pudong Foreign Language School is affiliated with SISU, Pudong New District, Shanghai, China
2) Cambridge Centre for International Research, Cambridge Science Park, Milton Road, Cambridge, UK; elenazhu2007@gmail.com
Mentor: Marco Vinicio Alban-Paccha

ABSTRACT: Heatstroke is a life-threatening condition that poses growing challenges under rising global temperatures. This study aimed to develop and evaluate machine learning models for predicting daily heatstroke incidence using nationwide health and meteorological data from Japan. Predictors included maximum, minimum, and average temperature, relative humidity, calculated heat index, calendar variables, and prefecture identifiers. Five modelling approaches were compared: Poisson regression, random forest, XGBoost, LightGBM, and LSTM. Model performance was assessed using mean squared error (MSE), root mean squared error (RMSE), and coefficient of determination (R^2). Among all models, LightGBM (MSE = 98.15, RMSE = 4.26, R^2 = 0.827) and XGBoost (MSE = 101.01, RMSE = 4.33, R^2 = 0.822) achieved better performance, outperforming both traditional Poisson regression (R^2 = 0.434) and neural network approaches (R^2 = 0.488). Feature importance analysis using SHapley Additive exPlanations (SHAP) confirmed maximum temperature as the strongest predictor. At the same time, the inclusion of prefecture identifiers substantially improved predictive power, underscoring the importance of regional heterogeneity. Excluding the prefecture reduced LightGBM performance to R^2 = 0.378. These findings highlight the potential of machine learning for early warning systems that integrate meteorological drivers and regional modifiers such as population. In the context of climate change and more frequent heatwaves, such adaptive models can support targeted public health strategies, optimize resource allocation, and reduce the health and economic burden of heatstroke.

KEYWORDS: Computational Biology and Bioinformatics, Computational Epidemiology, Heatstroke Prediction, Climate Change and Health, Machine Learning Models.

Introduction

Heatstroke is one of the most severe manifestations of heat-related illness, arising when the body's core temperature exceeds the compensatory capacity of thermoregulatory mechanisms. It can be triggered by environmental exposure to high temperatures, excessive physical exertion, or a combination of both, and typically presents as a spectrum of conditions ranging from mild heat exhaustion to potentially fatal heatstroke. The defining characteristics include a core body temperature exceeding 40 °C, central nervous system dysfunction (such as confusion, seizures, or coma), and progressive multi-organ failure. Without timely medical intervention, the case-fatality rate can exceed 50%, underscoring its urgency as a global public health issue.[1,2]

The burden of heatstroke is not limited to individual health outcomes. It has significant societal implications, particularly for vulnerable populations and high-risk occupational groups such as outdoor workers, athletes, emergency responders, and the elderly (Figure 1). Epidemiological studies consistently demonstrate a strong association between elevated ambient temperatures and an increase in hospital admissions, ambulance calls, and mortality attributable to heatstroke.[3] The Fire and Disaster Management Agency of Japan (FDMA) has reported a persistent upward trend in heatstroke-related emergency transports across the country. While annual fluctuations occur due to the presence or absence of extreme heat events and variations in rainy season length, the long-term trajectory reflects a clear increase in heat-related morbidity. This pattern is projected to worsen in the coming decades due to climate change, with intensifying heatwaves posing escalating risks to public health, occupational safety, and healthcare systems worldwide.

Beyond morbidity and mortality, heatstroke also carries a substantial economic burden. Hospitalizations for severe cases require intensive care, mechanical ventilation, and prolonged rehabilitation, placing a significant strain on healthcare systems during the peak summer months. Similar patterns have been observed internationally, where rising temperatures significantly increase hospital admissions and healthcare costs attributable to heat exposure.[4,5] The cost extends to productivity losses, as affected individuals may need extended recovery time away from work or experience long-term neurological and cardiovascular complications that reduce functional capacity. For employers, especially those in sectors reliant on outdoor labor, such as agriculture, construction, and transportation, heat-related illnesses contribute to absenteeism, reduced work performance, and occupational hazards. At a societal level, these direct and indirect costs combine to form a considerable economic impact, further justifying the urgent need for preventive strategies and predictive models.[6-8]

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.11

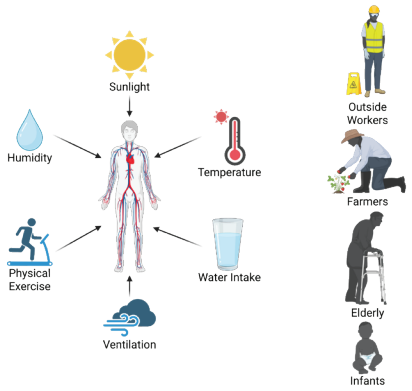


Figure 1: Environmental and behavioral factors contributing to heat stress (sunlight, temperature, humidity, sports, ventilation, and water intake) and the vulnerable groups most affected, including outdoor workers, farmers, the elderly, and infants.

Several studies illustrate this paradigm shift. Ogata *et al.* developed a machine learning model that combined multivariate meteorological and demographic data from 16 Japanese cities, achieving superior predictive accuracy over traditional Wet-Bulb Globe Temperature (WBGT)-based approaches for 12-hour incidence forecasting.[9] Xu *et al.* conducted a comprehensive comparison of machine learning algorithms, including random forest and gradient boosting methods, in a southern Chinese city and demonstrated that ensemble models significantly outperformed generalized linear models, particularly in capturing nonlinearities.[10] Shimazaki *et al.* advanced the field further by incorporating wearable sensor data (including predicted heat strain tolerance, heart rate, and activity level) into a supervised machine learning model that achieved 85.2 % predictive accuracy in occupational settings.[11] Meanwhile, Takada *et al.* demonstrated that simulated core body temperature could serve as a robust predictor of ambulance calls in Japanese metropolitan areas, performing comparably to environmental temperature-based models but with fewer input parameters.[12] Collectively, these studies highlight the growing utility of machine learning for predictive modelling of heat-related illness.

However, current literature has two key limitations. First, most studies have been restricted to single cities or a limited set of urban areas, thereby limiting the generalizability of their findings. Second, while meteorological and physiological features have been studied extensively, the potential contribution of geographical heterogeneity (e.g., regional socio-economic, demographic, and environmental differences) remains underexplored. This gap is particularly salient in a country like Japan, where diverse climatic zones, population structures, and levels of urbanization create region-specific vulnerabilities.

The present study addresses these gaps by conducting a nationwide analysis of heatstroke incidence across Japan using five years of health and meteorological data (2015–2019, June to September). By systematically comparing multiple machine learning approaches, including Poisson regression, random forest, XGBoost, LightGBM, and LSTM, the study evaluates predictive performance, identifies key features influencing outcomes, and assesses the role of regional variation in shaping heatstroke risk. In particular, the study applies SHapley Additive exPlanations (SHAP) to quantify feature importance, enabling interpretable insights into the drivers of heatstroke incidence.

Through this work, we aim to provide robust evidence for the integration of machine learning models into early warning systems for heat-related illness. By highlighting both the meteorological and regional determinants of heatstroke, this study contributes to the development of tailored public health strategies that account for geographical heterogeneity, ultimately supporting more effective risk prediction, prevention, and resource allocation in the face of escalating climate change.

Methods

Iterative machine learning workflow:

This study employs an iterative machine learning workflow aimed at systematically enhancing model performance through continuous refinement (Figure 2). The process commences with data preparation, which encompasses data integration and preprocessing. Following this, informative features are constructed via feature engineering, and models are trained using the training set. The trained models are subsequently evaluated on a test set. Should the performance fail to meet predefined criteria, feedback loops facilitate a return to earlier stages for further refinement. This "training-evaluation-feedback" cycle continues iterating until the model fulfills the established performance standards. Finally, the trained model is evaluated on a held-out test set to assess predictive performance and generalizability.

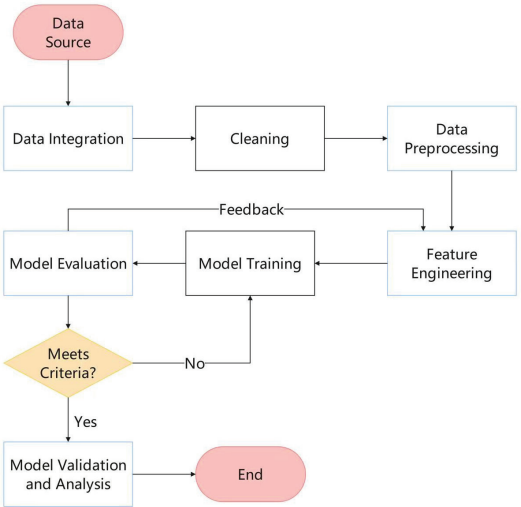


Figure 2: An iterative machine learning workflow guides the model development process, beginning with data sourcing and integration, moving systematically to cleaning, preprocessing, and feature engineering, and concluding with the model training phase. The model is then evaluated against predefined criteria, with feedback loops enabling cyclical refinement of data, features, or model parameters until performance standards are met.

Data sources:

This study utilized a publicly available dataset from Zenodo: *Datasets for quantifying association between short-term exposure to maximum temperature and heatstroke-related ambulance dispatches in Japan: a time-stratified case-crossover design* (https://zenodo.org/records/14560225), supplemented with data obtained from the official website of the Japan Meteorological Agency (JMA). The public dataset included daily records of

heatstroke-induced ambulance dispatches (HSAD) and meteorological data across Japan's 47 prefectures from 2015 to 2019, containing data of the summer months, June to September, when heatstroke incidence peaks. To enhance data completeness, daily minimum temperature and average temperature records were retrieved from the JMA website and integrated into the dataset. All data were sourced from official Japanese government channels, ensuring high reliability. The final dataset comprises 11 features and 28,200 records.

Study variables:

The outcome variable was the daily number of heatstroke cases per prefecture. Predictor variables included meteorological and temporal features:

- Daily maximum temperature (Tempmax, ˚C)
- Daily minimum temperature (Tempmin, ˚C)
- Daily average temperature (Tempave, ˚C)
- Average relative humidity (Rhumave, %)
- Calculated heat index (HeatIndex, composite variable of temperature and humidity)
- Calendar information: year, month, day of month, day of week (Dow), and holiday status (binary indicator)
- Prefecture identifier (categorical variable capturing regional variation)

To improve the predictive performance of the model, the heat index was added, a commonly used composite meteorological indicator. The Heat Index, also referred to as apparent temperature, is a meteorological measure that combines ambient temperature and relative humidity to reflect perceived thermal discomfort. Its primary function is to quantify the level of heat stress experienced by individuals under hot and humid conditions. High humidity impedes sweat evaporation, reduces the body's cooling efficiency, and thereby elevates the perceived temperature beyond the actual air temperature. As a result, the Heat Index is widely employed in public health, weather advisories, and occupational safety to assess heat-related illness risks, develop preventive strategies, and issue

$$HI = c_1 + c_2 \cdot T + c_3 \cdot RH + c_4 \cdot T \cdot RH + c_5 \cdot T^2 + c_6 \cdot RH^2 + c_7 \cdot T^2 \cdot RH + c_8 \cdot T \cdot RH^2 + c_9 \cdot T^2 \cdot RH^2$$

The Heat Index (HI) was calculated using the standard nine-coefficient Rothfusz regression equation adopted by the U.S. National Weather Service. In this equation, HI denotes the Heat Index, T represents air temperature expressed in degrees Fahrenheit (˚F), RH represents relative humidity (%), and the regression coefficients are defined as:

$C_1 = -42.379$,
$C_2 = 2.04901523$,
$C_3 = 10.14333127$,
$C_4 = -0.22475541$,
$C_5 = -6.83783 \times 10^{-3}$,
$C_6 = -5.481717 \times 10^{-2}$,
$C_7 = 1.22874 \times 10^{-3}$,
$C_8 = 8.5282 \times 10^{-4}$,
$C_9 = -1.99 \times 10^{-6}$.

Following the original National Weather Service formulation, temperature values were converted from degrees Celsius to degrees Fahrenheit before applying the regression equation.

Feature selection and explainability

To assess the relative importance of meteorological and temporal features, SHapley Additive exPlanations (SHAP) were applied. SHAP is a game-theoretic framework that attributes the contribution of each feature to the prediction outcome. A random forest regression model was used as the baseline for SHAP computation, chosen for its robustness in modelling nonlinear relationships and for the efficiency of the TreeExplainer algorithm within the SHAP library. The random forest was configured with 50 trees and a maximum depth of 5 to balance computational cost and prevent overfitting.

Data splitting strategy:

For the primary model comparison, the dataset was randomly divided into training and testing subsets. All models were trained using the same predictor variables and evaluated on identical train-test splits to ensure fair comparison across algorithms. To further assess robustness and address potential concerns regarding temporal dependence, an additional temporal hold-out validation was conducted. In this analysis, data from 2015–2018 were used for model training, while data from 2019 were reserved exclusively for testing.

Model development

Five modelling approaches were evaluated:

1. Poisson regression, a count model traditionally used in epidemiology for rare events.
2. Random forest regression is an ensemble of decision trees trained using bootstrap sampling and feature bagging.
3. Extreme Gradient Boosting (XGBoost) is a gradient boosting framework optimized for efficiency and performance.
4. Light Gradient Boosting Machine (LightGBM) is a gradient boosting framework using histogram-based algorithms for faster training.
5. Long Short-Term Memory (LSTM) neural network, a recurrent neural network architecture designed to capture sequential dependencies in time-series data.

All models were trained on the same input features and evaluated using identical training-test splits to ensure comparability. Hyperparameters for ensemble models were optimized using grid search. For LightGBM, the final configuration included a maximum of 31 leaves per tree, a learning rate of 0.05, 200 boosting iterations, and a random seed set to 42.

Model evaluation:

Model performance was evaluated using three standard metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the Coefficient of Determination (R^2):

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

$$RMSE = \sqrt{MSE}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)^2}$$

where n denotes the number of observations, y_i represents the observed heatstroke incidence for observation i, $\hat{y}_i$ represents

DOI: 10.36838/IJHSR820.11

the corresponding model prediction, and $\bar{y}$ denotes the mean observed heatstroke incidence across all observations.

These metrics collectively capture the accuracy (MSE), interpretability (RMSE, in the same units as the outcome variable), and explanatory power (R^2) of the models. Comparisons across models were made to determine the most effective algorithm for predicting heatstroke incidence.

Ethical considerations:

All data used in this study were aggregated, anonymized, and publicly available. As no individual-level or identifiable health data were used, ethical approval was not required.

Results and Discussion

Descriptive analysis and observable patterns:s

Between 2015 and 2019, a total of five years of nationwide data were analyzed for the summer months of June through September. Across all prefectures, heatstroke incidence showed marked temporal and spatial variation, with clear peaks during extreme heat events. Descriptive statistics indicated that meteorological variables, particularly temperature and humidity, were strongly associated with heatstroke incidence, whereas calendar-related features such as holiday status, year, and day of week had comparatively limited influence.

The calculated heat index demonstrated the strongest association with heatstroke counts among the meteorological variables, highlighting its utility as a composite indicator of thermal stress. Correlation analysis revealed coefficients of 0.48 for daily maximum temperature, 0.41 for minimum temperature, 0.47 for average temperature, and 0.51 for heat index, all with p-values ≤0.001 (Figure 3). These findings confirm that rising ambient temperatures and higher heat index values substantially increase the probability of heatstroke events. In contrast, variables such as year, month, day, and holiday status exhibited correlation coefficients below 0.1, suggesting their contribution was minor compared to direct meteorological drivers.

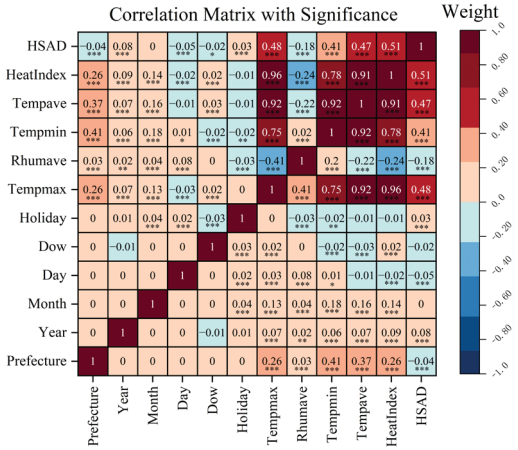


Figure 3: Correlation matrix with significance levels for meteorological, temporal, and regional features in relation to heatstroke admissions (HSAD). Strong positive correlations are observed between temperature variables and heat index, while weaker or negligible associations are seen for calendar-related features (year, month, day, holiday, and day of week).

Feature importance analysis:

To quantify the contribution of individual features, SHapley Additive exPlanations (SHAP) values were computed using a random forest model as the reference. The ranking of feature importance was: daily maximum temperature (Tempmax), prefecture identifier, daily average temperature (Tempave), heat index, daily minimum temperature (Tempmin), month, day, average relative humidity (Rhumave), year, and day of week (Dow). These results were used to rationalize and refine our feature set. Based on the SHAP-derived importance ranking, the study selected the six most influential predictors—daily maximum temperature (Tempmax), prefecture identifier, daily average temperature (Tempave), heat index, daily minimum temperature (Tempmin), and month—for the final model, ensuring it focuses on the dominant risk factors (Figure 4).

These results underscore two key findings. First, maximum temperature remains the single most powerful meteorological determinant of heatstroke incidence. Second, the inclusion of the prefecture identifier significantly improves model performance, pointing to strong regional heterogeneity. The prefecture variable effectively encodes differences in population structure, socio-economic conditions, urbanization level, and healthcare reporting standards, all of which modify the baseline vulnerability to heat-related illness.

Although relative humidity exhibited relatively low SHAP importance, this does not imply that humidity is unimportant for heatstroke risk. A likely explanation is that humidity information is already incorporated into the Heat Index variable. Because the heat index is calculated as a function of both temperature and relative humidity, substantial collinearity exists between these predictors. Consequently, part of the predictive contribution of humidity is captured by the Heat Index, reducing its standalone importance in the SHAP ranking.

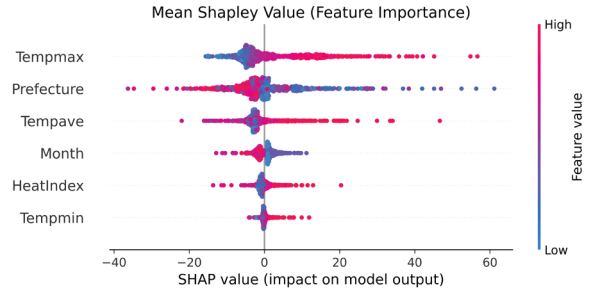


Figure 4: SHAP (SHapley Additive exPlanations) summary plot showing the relative importance of features in predicting daily heatstroke incidence. Maximum temperature (Tempmax) and prefecture identifier are the most influential predictors, followed by average temperature, heat index, and minimum temperature, while calendar-related features contribute minimally. The significant impact of the prefecture identifier suggests that localized factors, such as urban density or regional adaptation policies, critically modulate heatstroke risk beyond mere temperature. Consequently, this finding not only validates the model's design but also strongly advocates for the development of region-specific public health interventions.

Model comparison:

Five modelling approaches were evaluated: Poisson regression, random forest, XGBoost, LightGBM, and LSTM. Their performance was assessed using MSE, RMSE, and R^2 (Table 1).

Table 1: Predictive performance of machine learning and statistical models for daily heatstroke incidence, evaluated using mean squared error (MSE), root mean squared error (RMSE), and coefficient of determination (R^2). Notably, tree-based models, including XGBoost, Random Forest, and LightGBM, all achieved high explanatory power with R^2 values above 0.8, with LightGBM and XGBoost being the top performers. In contrast, both Poisson regression and LSTM demonstrated considerably lower predictive accuracy.

Model name	MSE	RMSE	R^2
Poisson regression	320.7344	7.4836	0.4340
XGBoost	101.0053	4.3259	0.8218
Random forest	108.6334	4.4663	0.8083
LSTM	290.4317	7.3702	0.4875
LightGBM	98.1498	4.2627	0.8268

Tree-based ensemble methods substantially outperformed the specific LSTM configuration evaluated in this study, as well as the traditional statistical model (Poisson regression). LightGBM and XGBoost achieved better overall performance. LightGBM achieved an MSE of 98.15 and an R^2 of 0.827, while XGBoost achieved an MSE of 101.01 and an R^2 of 0.822, representing a more than twofold improvement over Poisson regression. These results highlight the ability of gradient boosting algorithms to capture complex nonlinear interactions between meteorological factors and health outcomes. Because the performance differences among the top-performing models were relatively small, additional stability analyses were conducted using repeated random train–test splits to evaluate the robustness of these results.

Stability analysis across repeated random splits:

To evaluate the robustness of model performance across different random train–test partitions, the three best-performing machine-learning models were evaluated using 10 repeated random splits. The mean and standard deviation of R^2 values were calculated. LightGBM achieved a mean R^2 of 0.831 ± 0.017, XGBoost achieved 0.828 ± 0.020, and Random Forest achieved 0.809 ± 0.017. The difference between LightGBM and XGBoost was substantially smaller than the observed run-to-run variability, indicating that the two algorithms exhibited comparable predictive performance. Both models consistently outperformed Random Forest.

Table 2: Mean and standard deviation (SD) of R^2 values calculated from 10 repeated random train–test splits using the same sampling strategy as the primary analysis.

Model name	Mean R^2	SD
LightGBM	0.831	0.017
XGBoost	0.828	0.020
Random forest	0.809	0.017

Temporal hold-out validation:

To evaluate the robustness of the findings under a more stringent forecasting scenario, an additional temporal hold-out validation was performed using data from 2015–2018 for training and data from 2019 for testing.

The results demonstrated that the overall ranking of models remained largely unchanged. LightGBM and XGBoost achieved the highest predictive performance (R^2 = 0.769), followed closely by Random Forest (R^2 = 0.761). Poisson Regression and LSTM yielded substantially lower predictive accuracy, with R^2 values of 0.433 and 0.415, respectively. Although these values were modestly lower than those obtained using the random train-test split, the relative performance of the models remained consistent, indicating that the stronger performance of the tree-based ensemble models relative to the LSTM configuration evaluated in this study was robust to the choice of validation strategy.

Effect of regional heterogeneity:

The decisive role of regional characteristics was further examined through comparative experiments. When the prefecture feature was excluded, LightGBM performance dropped significantly (R^2 = 0. 378). With the prefecture feature included, performance improved dramatically to R^2 = 0.827. This pattern was consistent across all models tested, demonstrating that regional differences account for a substantial proportion of the variance in heatstroke incidence (Figure 5).

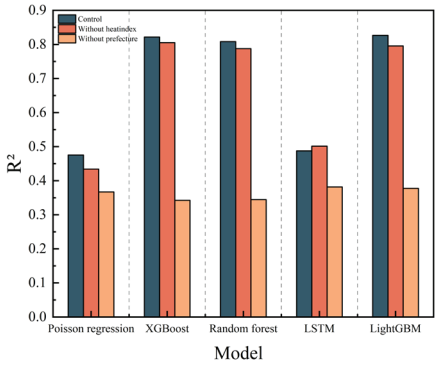


Figure 5: Comparison of predictive performance (R^2 values) across five models (Poisson regression, XGBoost, Random Forest, LSTM, and LightGBM) under three conditions: full model (control), without heat index, and without prefecture. The inclusion of prefecture notably improves accuracy across all models. The results reveal that the exclusion of the prefecture identifier significantly degraded model accuracy, whereas removing the heat index led to only a marginal decline.

Population adjustment analysis:

To examine whether the predictive contribution of the prefecture identifier primarily reflected differences in population size, an additional analysis was conducted by incorporating prefecture population as an explicit predictor.

The model, including meteorological variables and population, but excluding the prefecture identifier, achieved an R^2 of 0.850. This performance was comparable to that of the model including both population and prefecture information (R^2 = 0.849), suggesting that population size accounted for much of the predictive value associated with prefecture-level differences. These findings indicate that a substantial proportion of the regional variation observed in heatstroke case counts is related to population distribution across prefectures.

Interpretation of results:

LightGBM and XGBoost provided the most accurate predictions among the evaluated models and demonstrated strong robustness under both random train-test splitting and temporal hold-out validation. The models confirmed that meteorological drivers (maximum, minimum, and average temperature, as well as heat index) were the primary predictors, while regional differences, largely reflected by population distribution, provided additional predictive information.

DOI: 10.36838/IJHSR820.11

This dual influence of weather and region implies that national-level warning systems may be insufficient if they do not account for localized vulnerabilities. For example, prefectures with higher elderly populations or stronger urban heat island effects may experience disproportionately greater impacts under identical meteorological conditions.

Implications:

The findings of this study highlight several important implications for the prediction and prevention of heatstroke. The strong predictive performance achieved by the machine-learning models suggests that routinely available meteorological information can be effectively incorporated into operational early-warning systems and public-health decision support. Such approaches may help identify periods of elevated risk and facilitate timely interventions to reduce heat-related health burdens.

Equally significant is the role of regional heterogeneity. The inclusion of the prefecture feature markedly improved predictive performance across all models, suggesting that heatstroke risk cannot be fully explained by meteorology alone. While prefecture-level information was initially identified as one of the most influential predictors, additional analyses suggest that a substantial portion of this effect is attributable to population size. Larger prefectures such as Tokyo and Osaka naturally record more heatstroke cases because more individuals are exposed to heat-related risks. Therefore, caution is warranted when interpreting the prefecture identifier as a direct proxy for vulnerability, urban heat-island intensity, or socio-economic characteristics. Future studies incorporating demographic, socioeconomic, and environmental indicators may help disentangle these mechanisms.

These results gain additional importance in the context of climate change. Rising global temperatures and the increasing frequency, intensity, and duration of heatwaves are expected to exacerbate the health burden of heatstroke worldwide. Projections for Japan and other temperate countries suggest that populations with little historical exposure to extreme heat will face growing risks in the coming decades.[13, 14] This challenge will be compounded by global economic impacts, where labor productivity losses and healthcare demand are projected to escalate across multiple sectors.[15, 16] Under these circumstances, adaptive and flexible prediction models will be essential. Machine learning frameworks, capable of integrating new environmental parameters and adjusting to shifting climate baselines, provide a pathway to resilient early warning systems.

Together, the evidence indicates that accurate heatstroke prediction requires the integration of both meteorological drivers and regional modifiers, embedded within a forward-looking framework that accounts for climate change. By incorporating these elements, machine learning models can underpin dynamic early warning systems that guide healthcare resource allocation, inform government policy, and deliver targeted interventions to protect vulnerable populations in an era of increasing thermal stress.

Future research directions:

Although substantial regional heterogeneity in heatstroke risk was observed, the drivers of these spatial disparities require further investigation. Future studies should integrate demographic, socioeconomic, urbanization, and climate projection data to develop region-specific prediction models and assess future heatstroke burden. External validation is also needed to facilitate implementation in local heat-health early warning systems. Furthermore, while this study evaluated established machine learning algorithms, future research could explore modeling frameworks that explicitly account for spatial heterogeneity and regional interactions to further improve predictive performance and interpretability.

■ Conclusion

This study systematically analyzed the relationship between meteorological conditions and heatstroke incidence across Japan, using five years of nationwide data. By comparing multiple modelling approaches, we demonstrated that LightGBM and XGBoost achieved the strongest predictive performance for nationwide heatstroke forecasting using meteorological and regional information. The findings confirm that machine-learning approaches can effectively capture the complex relationships underlying heatstroke occurrence and provide a reliable basis for predictive modelling. These findings emphasize that both environmental exposure and local socio-geographical factors must be considered to achieve accurate and actionable predictions.

The results carry important implications for public health. By capturing complex nonlinear interactions, machine learning models provide a reliable foundation for early warning systems capable of anticipating spikes in heatstroke cases. More importantly, the strong effect of regional variation underscores the necessity of tailoring preventive strategies to local conditions rather than relying solely on national averages. Such approaches can enable healthcare providers and government agencies to allocate resources more efficiently and deliver targeted guidance to high-risk populations.

Looking forward, the importance of accurate prediction systems will only grow under the pressures of climate change. Rising global temperatures and more frequent heatwaves will amplify the health burden of heat-related illness, particularly in populations with limited adaptive capacity. Prediction models must therefore not only integrate meteorological and socio-demographic data but also remain adaptable to shifting climatic baselines. The development of dynamic, multi-scale systems that combine environmental monitoring, demographic insights, and machine learning prediction represents a critical step toward building climate-resilient public health infrastructure.

In conclusion, this study provides evidence that machine learning, when combined with meteorological and regional features, can serve as a powerful tool for forecasting heatstroke risk. By embedding these insights into flexible, climate-aware early warning platforms, it is possible to mitigate the health impacts of extreme heat, protect vulnerable populations, and

reduce the societal costs of climate-driven thermal stress, which are expected to rise sharply under future warming scenarios.[14-17]

Acknowledgments

The author wishes to sincerely thank Dr Marco Alban-Paccha for his mentorship, encouragement, and invaluable guidance throughout this research journey. His support not only strengthened the project but also deepened the author's passion for engineering solutions to real-world challenges.

References

1. Armed Forces Health Surveillance Division. Update: Heat illness, active component, U.S. Armed Forces, 2021. *MSMR* 2022, *29* (5), 2–8. https://health.mil/Reference- Center/Publications/2022/05/01/MSMR-Vol-29-No-5.
2. Song, Q.; Mao, H. D.; Liu, S. Y. Definition and classification of heat illness. *Med. J. Chin. People's Liberation Army* 2019, *44* (7), 541–545. https://doi.org/10.11855/j.issn.0577-7402.2019.07.01.
3. Professional Committee of Critical Care Medicine of the Chinese PLA. Expert consensus on standardized diagnosis and treatment for heat stroke (draft). *Med. J. Chin. People's Liberation Army* 2015, *40* (1), 1–7.https://doi.org/10.11855/j.issn.0577-7402.2015.01.01.
4. Liao, S., Pan, W., Wen, L., Chen, R., Pan, D., Wang, R., Hu, C., Duan, H., Weng, H., Tian, C., Kong, W. Temperature-related hospitalization burden under climate change. *Nature*. 2025, *644*, 960–968. https://doi.org/10.1038/s41586-025-09352-w.
5. Tong, M.X., Wondmagegn, B.Y., Williams, S., Hansen, A., Dear, K., Pisaniello, D., Xiang, J., Xiao, J., Jian, L., Scalley, B., Nitschke, M. Hospital healthcare costs attributable to heat and future projections. *Adv. Clim. Change Res*. 2021, *12* (2), 174–182.https://doi.org/10.1016/j.accre.2021.07.008.
6. Sapari, H., Selamat, M.I., Isa, M.R., Ismail, R., and Mahiyuddin, W.R.W. The Impact of Heat Waves on Health Care Services in Low- or Middle-Income Countries: Protocol for a Systematic Review. *JMIR Res Protoc*. 2023, *20*. https://pmc.ncbi.nlm.nih.gov/articles/PMC10616749.
7. Chen, F., Zhang, X., and Chen, Z. Behind climate change: Extreme heat and health cost. *Structural Change and Economic Dynamics*. 2023, *64*, 101-110.https://doi.org/10.1016/j.strueco.2022.12.007.
8. Boudreault, J.; Campagna, C.; Lavigne, E.; Chebana, F. Projecting the overall heat-related health burden and associated economic costs in a climate change context in Quebec, Canada. *Sci. Total Environ*. 2025, *958*, 178022. https://doi.org/10.1016/j.scitotenv.2024.178022.
9. Ogata, S.; Takegami, M.; Ozaki, T.; Nakashima, T.; Onozuka, D.; Murata, S.; Nakaoku, Y.; Suzuki, K.; Hagihara, A.; Noguchi, T.; Iihara, K.; Kitazume, K.; Morioka, T.; Yamazaki, S.; Yoshida, T.; Yamagata, Y.; Nishimura, K. Heatstroke predictions by machine learning, weather information, and an all- population registry for 12-hour heatstroke alerts. *Nat. Commun*. 2021, *12*, 4575. https://doi.org/10.1038/s41467-021-24823-0.
10. Xu, H.; Guo, S.; Shi, X.; Wu, Y.; Pan, J.; Gao, H.; Tang, Y.; Han, A. Machine learning-based analysis and prediction of meteorological factors and urban heatstroke diseases. *Front. Public Health* 2024, *12*, 1420608.https://doi.org/10.3389/fpubh.2024.1420608.
11. Shimazaki, T.; Anzai, D.; Watanabe, K.; Nakajima, A.; Fukuda, M.; Ata, S. Heat stroke prevention in hot specific occupational environment enhanced by supervised machine learning with personalized vital signs. *Sensors* 2022, *22* (1), 395.
12. Takada, A.; Kodera, S.; Suzuki, K.; Nemoto, M.; Egawa, R.; Takizawa, H.; Hirata, A. Estimation of the number of heat illness patients in eight metropolitan prefectures of Japan: Correlation with ambient temperature and computed thermophysical responses. *Front. Public Health* 2023, *11*, 1061135. https://doi.org/10.3389/fpubh.2023.1061135.
13. Anderson, G. B.; Bell, M. L.; Peng, R. D. Methods to calculate the heat index as an exposure metric in environmental health research. *Environ Health Perspect*. 2013, 121:1111–9. https://doi.org/10.1289/ehp.1206273.
14. Oyama, T.; Takakura, J.; Ishizaki, N. N.; Oka, K.; Honda, Y.; Masago, Y.; Hijioka, Y. Nationwide high- resolution heat risk projections and intervention cost analysis for the elderly in Japan under climate and demographic changes. SSRN preprint 2025. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5386992.
15. Ueta, H., Kodera, S., Sugimoto, S., and Hirata, A. Projection of future heat-related morbidity in three prefectures of Japan under 2 °C global warming scenarios. *Environ. Res*. 2024, *247*, 118202. https://doi.org/10.1016/j.envres.2024.118202.
16. Sun, Y., Zhu, S., Wang, D., Duan, J., Lu, H., Yin, H., Tan, C., Zhang, L., Zhao, M., Cai, W., and Wang, Y. Global supply chains amplify economic costs of future extreme heat risk. *Nature* 2024, *627*, 123– 129.https://doi.org/10.1038/s41586-024-07147-z.
17. Falchetta, G., De Cian, E., Sue Wing, I., and Carr, D. Global projections of heat exposure of older adults. *Nat. Commun*. 2024, *15*, 47197.https://doi.org/10.1038/s41467-024-47197-5.

Author

Jingyao Zhu is a 12th-grade student from Pudong Foreign Language School affiliated with SISU, in Shanghai, China, who is passionate about engineering for social good. She developed LifeTag, a smart garment with sensors to monitor sanitation workers' health under extreme heat. Combining physics knowledge, self-taught programming, and innovation contest experience, she plans to pursue engineering at university, applying data science and machine learning to global health challenges.

DOI: 10.36838/IJHSR820.11

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

When Social Proof Becomes Demand: Price Resistance and Consumer Behavior in Viral Food Trends

Elisa Lei

Adlai E Stevenson, 1 Stevenson Dr, Lincolnshire, IL, 60069, Lake; elisa.lei0102@gmail.com

ABSTRACT: Recently, certain foods like Dubai chocolate have exploded in popularity across social media platforms, despite these viral products being more expensive than normal products. But what exactly is the driving force behind participating in these viral food trends? This study examines whether social proof, defined as trying a viral food because others posted about it, talked about it, or recommended it, reduces price resistance in viral food markets. This study is grounded in theories such as informational cascades, social influence, media richness, and fear of missing out, along with surveying 271 students and adults in the Chicago area in a short survey with questions involving social media, price resistance, and what influenced participation in food trends. The survey data were then put into logistic regression models to determine significance and correlation. The analysis supports that social proof is strongly associated with lowering consumers' price resistance. However, social proof was found not to be a statistically significant predictor of repeat purchasing or posting and sharing behavior. Instead, purchase frequency was found to be the strongest predictor of both repeat purchasing and sharing, suggesting that long-term demand, unlike short-term demand, which depends on social proof, depends more on consumer habit and experience. Overall, the findings show that social proof is most powerful in generating initial demand for viral food products, but its effects weaken after the first purchase.

KEYWORDS: Economics, Behavioral Economics, Consumer Behavior, Social Proof, Viral Food Trends.

Introduction

The year 2025 saw the rise of certain viral food trends such as matcha, Dubai chocolate, Crumbl cookies, and protein snacks. More recently, even simple products like Dot cakes, which are just plain cakes with sprinkles, are exploding online. These foods became so widespread that they were and still are widely circulated across platforms such as TikTok, Instagram, and YouTube, and businesses and consumers alike were trying their best to participate in these trends. However, these food items are usually priced higher than items that aren't viral but otherwise similar in taste. For example, Dubai chocolate generally costs between $15 and $30, whereas a normal Hershey's bar costs just $1.50. Crumbl cookies can cost up to almost $7 a cookie, compared to one of their competitors, Insomnia Cookies, which only charges $3.50 per cookie. Yet, despite these high prices, the consumer demand does not seem to be slowing down. Looking through an economics lens, the high demand and widespread adoption of these foods reflect how social media is able to influence consumer demand. Consumers are not just influenced by taste but also by visual appeal and peer influence.

This study focuses on social proof as a central driver of consumer behavior in viral food markets. Social proof is defined here as trying a viral food because other people posted about it, talked about it, or recommended it. The main research question is whether social proof reduces consumers' price resistance in viral food markets. To answer this question, surveys were sent out to high school students and adults in the Chicago area, and their responses were collected, measuring exposure to viral food trends, social proof, perceived overpricing, repeat purchasing, willingness to pay, and social media influence.

This study is grounded in several theories combining ideas like consumer utility theory, informational cascades, social influence, media richness, scarcity, prospect theory, and fear of missing out (FOMO). In a simple utility framework, a consumer purchases when the perceived utility of buying exceeds the perceived cost:

$$\text{ConsumerUtility}_i = \text{TasteUtility}_i + \text{SocialParticipationUtility}_i + \text{Scarcity/FOMOUtility}_i + \text{InformationSignal}_i - \text{PriceCost}_i$$

Social proof can increase the nonprice utility of a purchase by making a product appear more popular, trustworthy, urgent, and socially meaningful. Herd behavior and informational cascade theories explain why consumers may rationally imitate others when product quality is uncertain. Because cascades can form from a small number of early signals, they also help explain why viral food products can become popular quickly and then lose attention just as quickly.[1,2]

Social influence theory adds that consumers may respond to others through compliance, identification, and internalization. In viral food markets, internalization occurs when consumers treat reviews, ratings, and influencer recommendations as useful information; identification occurs when consumers try foods to feel connected to a broader trend community; and compliance occurs when consumers respond to perceived social pressure or urgency.[3] Media richness theory further helps explain why short videos, photos, and influencer posts are powerful for food products: these formats can communicate taste, texture, atmosphere, and emotional reaction more effectively than text-only information.[4]

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.12

Empirical studies provide added support for these theories. Prior research on digital engagement suggests that social media advertisements, electronic word of mouth (e-WOM), and social media influencers can increase consumer responsiveness in the food and beverage industry.[5] Viral marketing research also suggests that useful, credible, and informative content can strengthen consumer attitudes and purchase intentions.[9] These mechanisms matter because they show how consumers may use posts, comments, reviews, and ratings as external signals of product value, because they are initially unable to fully evaluate taste, quality, or experience themselves.

Viral food purchases may also provide additional utility from participation. Consumers may buy a product to post about it, discuss it with peers, or feel included in a trend. Research on coffee shops, for example, shows how businesses use social media to build brand awareness, trust, and loyalty, while using their visually appealing spaces to encourage customers to visit, take photos, and post about it online.[6] As these posts gain views, they increase the visibility of the product and the experience associated with it, making more people want to participate. Social media can also intensify FOMO by making other people's purchases and experiences highly visible, creating urgency and an increasing desire to participate. When a food product is framed as limited in time or quantity, consumers may interpret not buying as losing a desirable opportunity. Prospect theory helps explain this behavior because anticipated losses may feel more psychologically powerful than equivalent gains.[7,8]

This study contributes to the literature and previous studies by testing whether social proof is associated with lower price resistance in a specific consumer context: viral food markets. It also separates initial purchase from repeat purchase and posting behavior. The study tests three hypotheses: H1, social proof is positively associated with purchasing a viral food despite perceiving it as overpriced; H2, social proof is positively associated with repeat purchasing; and H3, social proof is positively associated with posting or sharing viral food content.

Methods

Using a cross-sectional survey design, two Google Form surveys were sent out, one to high school students and one to adults in the Long Grove/Chicago, Illinois, area. The combined consumer sample includes 271 respondents, consisting of 79 students and 192 adults. Due to a few respondents skipping specific questions, some variables have smaller observation counts. The survey was designed to measure whether social proof is associated with three consumer outcomes: purchase despite perceived overpricing, repeat purchase, and posting or sharing behavior.

The survey design was informed by three case examples of viral food trends: Dubai chocolate, matcha, and Crumbl. Dubai chocolate illustrates how online visibility and scarcity can transform a niche product into a high-demand viral item, with major retailers and food chains later responding to the trend.[10] Matcha illustrates a trend that combines visual appeal, health identity, and sustained market growth, shown through the long-lasting attention and continued product adoption by food businesses.[11,15] Crumbl illustrates how rotating flavors, distinctive packaging, and frequent social media reviews can turn limited-time flavors into recurring engagement, despite criticism that the product may be overhyped.[16,17]

Table 1 defines the consumer variables and reports descriptive statistics. For binary variables, the mean represents the proportion of respondents coded as 1. Variable coding was harmonized across student and adult survey files so that similar concepts with different question wording were represented by standardized binary or ordinal variables (Table 2).

Table 1: Consumer variable definitions and descriptive statistics. Descriptive statistics show that social proof and social media exposure were common in the sample, as 75.2% of valid respondents reported that they tried a viral food because others posted, talked about, or recommended it, and 68.3% first encountered viral food trends through social media. However, only 19.0% reported posting or sharing viral food content, suggesting that consumers were more likely to be influenced by viral food trends than to actively spread them.

Variable	Definition	Obs.	Mean	Std. Dev.	Min	Max
OverpricedPurchase	Whether the respondent purchased a viral food even though they believed it was overpriced.	269	0.595	0.492	0	1
SocialProof	Whether the respondent tried a viral food because others were posting about it, talking about it, or recommending it.	270	0.752	0.433	0	1
RepeatPurchase	Whether the respondent continued buying a viral food at least occasionally after the first purchase.	265	0.577	0.495	0	1
FOMO/Urgency	Whether the respondent said food trends feel limited-time, urgent, or time-sensitive.	266	0.549	0.499	0	1
FirstSocialExposure	Whether the respondent first learned about or saw viral food trends through a social media platform.	271	0.683	0.466	0	1
PostShare	Whether the respondent posted or shared a trending food on social media.	269	0.190	0.393	0	1
TasteReason	Whether taste was selected as a reason for continued purchasing.	271	0.878	0.328	0	1
Student	Whether the respondent belongs to the student survey sample.	271	0.292	0.455	0	1
SocialMediaInfluence	Self-reported level of social media influence on food-trend decisions, coded from 0 to 4.	270	1.348	1.004	0	4
PurchaseFrequency	Frequency of purchasing viral or trending foods, coded from 0 to 4.	270	1.548	0.860	0	4
MaximumWTP	Maximum willingness to pay for a viral food item, coded in dollars.	270	13.778	7.843	0	40

Note. For binary variables, the mean represents the proportion of respondents coded as 1. For example, the mean of SocialProof is 0.752, meaning approximately 75.2% of valid respondents reported trying a viral food because others posted, talked about, or recommended it. MaximumWTP is coded as follows: adult responses of 0 are coded as 0; "under $5" = 5; "~$10" = 10; "~$20" = 20; "~$30" = 30; student "no limit" responses are coded as 40.

Table 2: Variable construction and coding.

Variable	Coding
OverpricedPurchase	1 = purchased despite believing item was overpriced; 0 = did not
RepeatPurchase	1 = purchased again after first trial; 0 = did not
PostShare	1 = posted/shared viral food online; 0 = did not
SocialProof	1 = tried because others posted/recommended/talked about it; 0 = no
Student	1 = student respondent; 0 = adult respondent
PurchaseFrequency	0 (never) to 4 (very often)
MaximumWTP	0–40 scale from willingness-to-pay categories
SocialMediaInfluence	0 (not influential) to 4 (extremely influential)
FirstSocialExposure	1 = first learned via social media; 0 = otherwise

The primary model tests whether social proof reduces price resistance by increasing the probability that consumers purchase a viral food despite perceiving it as overpriced. Let $P_{op,i}$ represent the probability that respondent *i* purchased a viral food despite perceiving it as overpriced:

$$P_{op,i} = \Pr(\text{OverpricedPurchase}i=1)$$

The model is specified as:

$$\ln(P_{op,i} / (1-P_{op,i})) = \beta_0 + \beta_1 SocialProof_i + \beta_2 Student_i + \beta_3 PurchaseFrequency_i + \beta_4 MaximumWTP_i + \beta_5 SocialMediaInfluence_i + \beta_6 FirstSocialExposure_i + \varepsilon_i$$

The repeat purchase model examines whether the same viral demand factors that encourage initial purchase also sustain consumer demand over time. Let Prp,i represent the probability that respondent i purchases a viral food again after the initial trial:

$$P_{rp,i} = Pr(RepeatPurchasei=1)$$

The model is specified as:

$$\ln(P_{rp,i} / (1-P_{rp,i})) = \alpha_0 + \alpha_1 SocialProof_i + \alpha_2 Student_i + \alpha_3 PurchaseFrequency_i + \alpha_4 MaximumWTP_i + \alpha_5 SocialMediaInfluence_i + \alpha_6 FirstSocialExposure_i + u_i$$

The sharing and diffusion model evaluates whether participation in viral food trends extends into online posting or sharing behavior. $P_{ps,i}$ represent the probability that respondent i posted or shared viral food content:

$$P_{ps,i} = Pr(PostSharei=1)$$

The model is specified as:

$$\ln(P_{ps,i} / (1-P_{ps,i})) = \delta_0 + \delta_1 SocialProof_i + \delta_2 Student_i + \delta_3 PurchaseFrequency_i + \delta_4 SocialMediaInfluence_i + \delta_5 FirstSocialExposure_i + \varepsilon_i$$

In each model, the left-hand side represents the log odds of the corresponding outcome occurring. All three outcomes are binary. Therefore, the models are estimated as logistic regressions. Results are reported as coefficients, odds ratios, average marginal effects, and p-values. Because the survey is cross-sectional and self-reported, the models identify associations rather than causal effects.

Results and Discussion

Descriptive statistics:s

The survey results show that viral food behavior is common among both students and adults. Across the combined sample, many respondents reported exposure to viral food trends through social media, and 75.2% of valid respondents reported trying a viral food because others posted, talked about, or recommended it (Table 1). The main outcome variable, OverpricedPurchase, measures whether a respondent purchased a viral food even though the respondent believed it was overpriced. This variable is used as the study's measure of reduced price resistance. The secondary outcomes are RepeatPurchase, which measures whether respondents continued buying the product after the first trial, and PostShare, which measures whether respondents posted or shared viral food content on social media.

Price resistance:

Table 3 presents the logistic regression results for the price resistance model. The dependent variable is OverpricedPurchase, equal to 1 if the respondent purchased a viral food despite believing that it was overpriced.

Table 3: Logistic regression predicting purchase despite perceived overpricing. Social proof is strongly and significantly associated with purchasing despite perceived overpricing. Respondents exposed to social proof had about 3.09 times higher odds of buying an overpriced viral food, and social proof was associated with a 21.6 percentage-point increase in the probability of purchase.

Predictor	Coefficient	Odds Ratio	Marginal Effect	p-value
Social Proof	1.127	3.087	0.216	0.0015
Student	-0.151	0.860	-0.029	0.6575
Purchase Frequency	0.652	1.918	0.125	0.0015
Maximum WTP	0.013	1.013	0.003	0.5444
Social Media Influence	0.398	1.489	0.076	0.0311
First Social Exposure	-0.504	0.604	-0.097	0.1421

Note. N = 268. McFadden pseudo R^2 = 0.163. AUC = 0.756. Marginal effects are average marginal effects.

The results support H1. Social proof is statistically significant and positively associated with reduced price resistance (p = 0.0015, Table 3). After controlling for student status, purchase frequency, maximum willingness to pay, social media influence, and first exposure through social media, respondents who experienced social proof had approximately 3.09 times higher odds of purchasing a viral food despite perceiving it as overpriced (Table 3). The average marginal effect suggests that social proof is associated with a 21.6 percentage-point increase in the probability of buying despite perceived overpricing (Table 3). Purchase frequency is also statistically significant (p = 0.0015, Table 3), indicating that respondents who buy trending foods more often are more likely to purchase despite overpricing. Social media influence is positive and statistically significant as well (p = 0.0311, Table 3).

Student status, maximum willingness to pay, and first exposure through social media are not statistically significant in the full model (Table 3).

Repeat purchase:

Table 4 presents the logistic regression results for the repeat purchase model. The dependent variable is RepeatPurchase, equal to 1 if the respondent continued buying a viral food after the first purchase.

Table 4: Logistic regression predicting repeat purchase. Social proof does not significantly predict repeat purchasing after the first trial. Instead, purchase frequency is the strongest predictor, suggesting that continued demand depends more on existing buying habits than on social proof alone.

Predictor	Coefficient	Odds Ratio	Marginal Effect	p-value
Social Proof	0.251	1.286	0.054	0.4631
Student	0.034	1.035	0.007	0.9149
Purchase Frequency	0.694	2.003	0.150	0.0004
Maximum WTP	0.028	1.028	0.006	0.1823
Social Media Influence	0.105	1.111	0.023	0.5438
First Social Exposure	-0.286	0.751	-0.062	0.3765

Note. N = 264. McFadden pseudo R^2 = 0.091. AUC = 0.699. Marginal effects are average marginal effects.

The repeat purchase results do not support H2. Social proof has a positive coefficient, but it is not statistically significant (p = 0.4631, Table 4). Therefore, the data do not provide strong evidence that social proof alone predicts whether consumers continue buying a viral food after the initial trial. The strongest

DOI: 10.36838/IJHSR820.12

predictor of repeat purchase is PurchaseFrequency (p = 0.0004, Table 4). Each one-point increase in the purchase frequency scale is associated with approximately 2.00 times higher odds of repeat purchase and a 15.0 percentage-point increase in the probability of repeat purchase (Table 4). This suggests that repeat purchasing is more closely related to existing food-trend buying habits than to social proof alone. Student status, maximum willingness to pay, social media influence, and first exposure through social media are not statistically significant in the full model (Table 4).

Posting and diffusion:

Table 5 presents the sharing and diffusion model. The dependent variable is PostShare, equal to 1 if the respondent posted or shared a trending food on social media.

Table 5: Logistic regression predicting posting or sharing behavior. Social proof is positively related to posting or sharing behavior, but the relationship is not statistically significant. Purchase frequency is the only statistically significant predictor, suggesting that consumers who buy viral foods more often are also more likely to share them online.

Predictor	Coefficient	Odds Ratio	Marginal Effect	p-value
Social Proof	0.755	2.128	0.099	0.1958
Student	0.120	1.127	0.016	0.7461
Purchase Frequency	0.736	2.088	0.097	0.0012
Social Media Influence	0.246	1.278	0.032	0.2369
First Social Exposure	0.640	1.897	0.084	0.1946

Note. N = 268. McFadden pseudo R^2 = 0.137. AUC = 0.733. Marginal effects are average marginal effects.

The sharing model provides limited support for H3. Social proof is positively associated with posting or sharing behavior, with an odds ratio of 2.13, but the relationship is not statistically significant (p = 0.1958, Table 5); therefore, it cannot be concluded that differences in sharing are actually due to social proof. First exposure through social media and social media influence are also positive, but not statistically significant (Table 5). The only statistically significant predictor is PurchaseFrequency (p = 0.0012, Table 5). Because posting or sharing may happen after purchase, this model should be interpreted as a diffusion outcome rather than as evidence that sharing causes demand. The available data suggest that sharing behavior may describe how viral food trends spread, but the timing of the survey does not allow a stronger causal claim.

Interpretation:

Taken together, the findings support a two-stage interpretation of demand in viral food markets. Social proof appears to be the most powerful in the early stage by reducing price resistance and encouraging consumers to try viral foods even if they perceive them as overpriced (Table 3). This pattern is consistent with the idea that visibility, endorsement, and trend participation can temporarily increase perceived utility. In later stages, however, social proof weakens as a predictor. Repeat purchasing and posting behavior depend more on consumers' existing engagement with food trends, as measured by purchase frequency (Tables 4 and 5). Therefore, this also explains why some trends die out quickly, especially if consumers originally didn't purchase the product frequently.

These findings are consistent with the theoretical predictions discussed earlier, particularly informational cascade theory and social influence theory. The strong association between social proof and purchasing despite perceived overpricing (Table 3) suggests that social information and participation value may increase consumers' perceived utility enough to offset the higher price. The fact that social proof predicts initial but not repeat purchasing (Table 4) is also consistent with informational cascade theory. Cascades can be fragile because they only rely on a small number of earlier decisions rather than independent evidence of product quality. Once consumers gain firsthand experience, these external signals may become less influential. Prospect theory, particularly loss aversion, may further explain this behavior when consumers view missing a limited or popular product as a loss, while FOMO may intensify the urgency to participate. However, because the survey did not directly measure loss aversion or experimentally manipulate scarcity, the results should not be interpreted as a direct test of prospect theory.

These findings can also be used as insight and a foundation for food businesses and marketers who are looking to utilize social media to attract more attention. Businesses should use marketing and product strategies that make use of social proof rather than just introducing new items that fit the current trends in order to attract more customers. Utilizing things like influencers, customer content, and reviews can be useful in helping spark initial demand for a product, even if pricing is on the higher end (Table 3). However, to turn these customers into repeat ones requires more emphasis on product quality and value (Table 4). In other words, online hype may attract attention, but long-term demand depends on product experience.

Limitations and future research:

Several limitations should be considered when interpreting the results. For starters, all survey respondents came from the Chicago area, which means the findings really only represent this one specific area and do not provide the greatest generalizability to other regions or cultural contexts. Additionally, the data is self-reported, which risks biases and certain respondents not being one hundred percent truthful. The study is also cross-sectional, which means it cannot determine causation. Next, posting and sharing outcomes present a timing issue because sharing may occur after purchase. As a result, the sharing model should be interpreted cautiously. The models may also be affected by unobserved confounding factors. For example, income, general price sensitivity, impulsivity, personality traits, or a desire for social belonging could influence both a consumer's responsiveness to social proof and willingness to purchase an overpriced viral food. Because these characteristics were not measured, the estimated relationship between social proof and price resistance may partly reflect differences among respondents that are not captured by the included control variables. Finally, these findings may only reflect the specific products and social media environment at the time of the survey, because viral food trends tend to change quickly.

To combat these limitations and add to this study, future research should consider surveying larger samples with respondents from different areas, not just one. Future research could collect measures of income, baseline price sensitivity, impulsivity, and social conformity. Experimental studies that randomly vary the amount of social proof shown to participants could also provide stronger evidence about causality. In addition, to distinguish between short-term trial and long-term repeat purchasing, future research can consider using longitudinal designs that could track consumers before and after exposure to a viral food product. Another idea could be comparing differences between social proof on different platforms like TikTok, Instagram, and YouTube, and their effects on price resistance. Future studies can also compare differences between short-lived trends and more prevalent trends, such as matcha, to pinpoint why some trends fade quickly while others last longer.

Conclusion

This study investigated how social proof shapes consumer behavior in viral food markets, with a particular focus on whether social proof reduces price resistance. Through surveys and models, the study found that social proof is significantly associated with reduced price resistance, which helps explain why consumers who were exposed to social proof were more likely to purchase viral foods even though they were overpriced. This suggests that social influence can temporarily weaken traditional price sensitivity when a product feels popular, urgent, or socially valuable, explaining why demand for viral foods is so high even though they are more expensive.

However, at the same time, the findings also show that social proof has its limitations. In the full models, social proof was found not to be a significant factor when it comes to customers purchasing again or posting and sharing online afterward. Instead, purchase frequency was found to be the strongest predictor. These findings suggest that social proof may contribute to initial purchases of viral foods, but long-term demand depends more on consumer habits and the actual product experience, reflecting why trends sometimes die out after a few months. Still, these findings can provide valuable insight for businesses, reflecting how social media visibility may be useful for attracting attention and creating initial demand for a product, but long-term success still depends on whether the product can satisfy consumers after the hype.

AI Acknowledgments

During the preparation of this work, the author used ChatGPT to revise formatting and tables, improve language readability, reorganize sections, and align the manuscript with the International Journal of High School Research template requirements. After using this tool, the author reviewed and edited the content as needed and took full responsibility for the content of the publication.

Acknowledgments

Firstly, I would like to thank my teachers for their lessons on starting a research paper. Thanks to their guidance, I have a solid foundation of the steps I needed to take. I would also like to thank my school library for providing access to a research database that includes a wide range of peer-reviewed sources. This resource greatly supported my research and helped me build a stronger foundation of existing theories upon which my study was able to build upon. In addition, I am very grateful to all 271 survey respondents who took time out of their day to participate in my survey. Thanks to them, I was able to make valuable conclusions and continue my study. Finally, I would like to thank my parents for their encouragement, feedback, and helpful advice throughout the research process.

References

1. Banerjee, A. V. A Simple Model of Herd Behavior. The Quarterly Journal of Economics 1992, 107 (3), 797–817. https://doi.org/10.2307/2118364.
2. Bikhchandani, S.; Hirshleifer, D.; Welch, I. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. Journal of Political Economy 1992, 100 (5), 992–1026. https://doi.org/10.1086/261849.
3. Kelman, H. C. Compliance, Identification, and Internalization: Three Processes of Attitude Change. Journal of Conflict Resolution 1958, 2 (1), 51–60.
4. Daft, R. L.; Lengel, R. H. Organizational Information Requirements, Media Richness and Structural Design. Management Science 1986, 32 (5), 554–571.
5. Sarkis, N.; Jabbour Al Maalouf, N.; Al Geitany, S. The Power of Digital Engagement: Unveiling How Social Media Shapes Customer Responsiveness in the Food and Beverage Industry. Administrative Sciences 2025, 15 (7), 278. https://doi.org/10.3390/admsci15070278.
6. Wibisana Sudarta, I. B.; Mahyuni, L. P. Viral Marketing: Utilizing Digital Content Digitization to Enhance Brand Awareness, Brand Trust, and Customer Loyalty in Coffee Shops. Eduvest 2024, 4 (6), 4677–4688. https://doi.org/10.59188/eduvest.v4i6.1146.
7. You, Y. Social Media Sharing, Fear of Missing Out, and Impulse Purchase: A Conceptual Study. Advances in Consumer Research 2025, 2 (4).
8. Kahneman, D.; Tversky, A. Prospect Theory: An Analysis of Decision under Risk. Econometrica 1979, 47 (2), 263–291.
9. Singh, P.; *et al.* The Viral Effect: Unpacking the Influence of Viral Marketing Campaigns on Generation Z's Purchase Intentions. SAGE Open 2025, 15 (2). https://doi.org/10.1177/21582440251346110.
10. Rice, N. Why Everyone Is Selling Dubai Chocolate Bars. CNBC, May 31, 2025. https://www.cnbc.com/2025/05/31/walmart-shake-shack-trader-joes-dubai-chocolate.html (accessed Apr 26, 2026).
11. Lawton, M. Companies Look to Make Green Off of Matcha's Growing Popularity. Marketplace, Oct 29, 2024. https://www.marketplace.org/story/2024/10/29/companies-look-to-make-green-off-of-matchas-growing-popularity (accessed Apr 26, 2026).
12. Google Trends. Explore. Google Trends, 2024. https://trends.google.com/explore?date=now%201-d&geo=US&q=%2Fg%2F11q98yn3fh (accessed Apr 26, 2026).
13. Grand View Research. Matcha Market Size, Share | Global Industry Trends Report, 2018–2025. https://www.grandviewre-

DOI: 10.36838/IJHSR820.12

search.com/industry-analysis/matcha-market (accessed Apr 26, 2026).
14. Yousif, S. 5 Matcha Pop-Ups That Are Actually Worth the Trip. Observer, Dec 15, 2025. https://observer.com/list/best-matcha-pop-ups-united-states/ (accessed Apr 26, 2026).
15. Mactas, A. Jack in the Box Just Put Matcha on the Menu & Fans Are Responding in All Caps: "EPIC." Delish, Feb 24, 2026. https://www.delish.com/food-news/a70466864/jack-in-the-box-now-serving-matcha/ (accessed Apr 26, 2026).
16. Kubzansky, W. Wall Street Wants In on Crumbl's Oversized Sugary Cookies. Bloomberg.com, June 2025.
17. Haydon, H. Are Crumbl Cookies Becoming Overrated? The Captain's Log, 2024. https://cnuclog.com/3150/uncategorized/opinion/are-crumbl-cookies-becoming-overrated/.

Author

Elisa Lei is a high school student at Adlai E Stevenson. Her academic interests include behavioral economics, consumer behavior, food markets, and the intersection of economics and mathematics, particularly econometrics. She is especially interested in using data and economic theory to better understand how consumers make decisions and how markets respond to changing technology and social trends. Elisa's interest in economics has also extended beyond the classroom. She placed in the top 10 in Illinois in the National Economics Competition, an experience that strengthened her interest in applying economic concepts to real-world issues. She plans to major in economics in college and hopes to continue exploring topics related to econometrics, consumer behavior, and food-related markets.

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

The Electric Vehicle in the Automotive Industry - A Review of the Economic Barriers That Limit Its Large-Scale Adoption

Kevi Sahara Karume

Braeside High School, Muthangari Road, Lavington, Nairobi, P.O. Box 25578-00603, Kenya; kevikarume@gmail.com
Mentor: Marcin Kedziera

ABSTRACT: The electric vehicle is a sustainable solution designed to reduce the environmental impact caused by the automotive industry. It replaces the traditional combustion-engine vehicle with a drivetrain powered by battery technology, resulting in zero emissions at the point of use. Yet, as several studies have pointed out, the success of this mass-scale replacement has been minimal. This paper aimed to explain the cost barriers and implications that deter buyers from purchasing electric vehicles (EVs). Additionally, this paper considered whether the electric vehicle would be a beneficial solution for the current economy and automotive market. I utilized literature reviews on electric vehicle production systems and secondary data analysis of statistics to illustrate the effectiveness of EVs in the industry. The EV is nowhere near replacing the demand for the combustion-powered vehicle, due to its high initial price, low range, and lack of charging infrastructure. This is a barrier on the road to a sustainable future. Questioning the reasons for such failure will spark reviews of the weaknesses of the solution, leading to higher purchase rates of electric vehicles, depending on the nation and energy fuel mix.

KEYWORDS: Automotive Engineering, Electric Vehicle, Total Cost of Ownership, Internal Combustion Engine Vehicle, Automotive Industry.

Introduction

The automotive industry is responsible for 28% of the emissions released into the atmosphere.[1] The mainstream use of internal combustion engine vehicles (ICEVs) that rely on the burning of hydrocarbons to generate chemical energy results in high emission rates due to the release of carbon dioxide and other particulates.[2,3] Despite the major environmental implications of the internal combustion engine vehicle (ICEV), its demand increases despite the various technological advancements in alternative solutions that promote zero emissions.

The alternative solution discussed in this review is the electric vehicle (EV). The powertrain of this vehicle is fueled by electric energy, releasing negligible tank-to-wheel emissions, which is preferable for the environment.[4] This does not take into account charging effects on emissions.[5,6] Several governments have pledged a complete transition to EVs in their countries. These nations further enforce policies that promote the production and purchase of EVs as part of the international commitment towards a sustainable future.[7] However, these efforts have not led to large-scale growth in the EV industry. This raises questions as to why electric vehicles are not purchased at such a high rate.

Electric vehicles are significantly more efficient (from tank to wheel) than traditional gasoline-powered vehicles.[4] The EV translates 60% of the energy generated to the wheels in comparison to 20%-40% in the ICEV.[1] In the year 2019, the total number of EV cars globally was 1%.[7] Such low purchase rates of EVs are caused by several parameters outside the limits of sustainability. Several studies claim this is still a majorly unexplored territory[4,7] with few reports analyzing the factors contributing to the stunted growth of the industry.[1,4]

The paper is a narrative review with an exploratory secondary data analysis, which has the main objective of identifying and evaluating the reasons for the slow adoption rate of electric vehicles. Through an in-depth analysis of several total cost of ownership (TCO) reports and life cycle cost analysis (LCCA) studies, the paper gives insight into the factors that cause the automotive market to lean towards internal combustion engine vehicles. Additionally, we utilize an exploratory secondary data analysis of statistics on the rate of purchase and regional income budgets to explore the affordability gap between the two types of vehicles. Moreover, the study considers solutions to the issues suggested by other papers cited in this review. This research adds to existing literature by comparing economic factors with EV sales, giving a visual demonstration of ICEV and EV purchase rates, and elaborating on the weaknesses in the system that lead to the low purchase rates of EVs, which few papers focus on. Furthermore, the review will explore the external parameters and answer the important yet overlooked question: "Why is the EV unable to replace the demand for the ICEV?"

This paper is structured as follows. The first section covers the main limitations of the EV, leading to a low purchase rate. The second section of the paper covers the causes of the upfront price and high cost of production of the EV. The third section covers the methodologies commonly used by the papers cited in this review. The third section also includes a short secondary data analysis comparing the wealth distribution and the sales of EVs, showing the effect of price on market marginalization, and comparing the price and the range of several EVs, which would suggest the effect of increasing the battery capacity on the price of the EV. The fourth section covers the

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.13

solutions suggested to solve the named limitations that lead to the slow progress of the EV market.s

Main limitations of the EV

Initial cost of purchase:

As discussed in the introduction, the Electric Vehicle is considered to be significantly more efficient than the internal combustion engine vehicle. The electric vehicle is a better long-term solution due to a lower frequency of maintenance and reduced fuel costs, leading to a lower total cost of ownership (TCO), which is the sum of all the expenses of a product, in comparison to the ICEV.[4,8]

The Consumers' focus on the initial purchase price serves as a major limitation towards EV adoption.[9–11] As several studies indicate, the initial purchase costs of EVs are higher than those of ICEVs, leading to the marginalization of several economic groups with lower incomes, due to the solution exceeding their financial ability.[7,9,12]

According to a research study comparing the total cost of ownership of the EV in comparison to the ICEV, the main concern stated by consumers is whether the cost of a greener solution "can be recouped within a reasonable time of ownership."[8] This statement further emphasizes the claim that consumers are concerned as to whether such a high price is worth the benefits of the EV, further illustrating the uncertainty in the market caused by the high upfront costs of the vehicle.

In addition, due to the advanced technology applied to the EV drive-train, the cost of general maintenance and repairs of the EV is high in comparison to the ICEV.[9] This is another limitation that makes the market prefer ICEVs, as the complications that come with intricate and delicate electronics, paired with the uncertainty of the new technologies, further cause lower purchase rates.

It is of the essence that companies produce a car that is affordable to the masses, and is at a competitive price to the ICEV. Otherwise, a major part of the automotive market will be financially excluded from purchasing a greener solution. This is a compromise that contains several complications, as discussed in the fourth section of this paper.

Range and convenience of fueling the Electric Vehicle:

Other limitations that concern consumers are the range of the EV and the availability of the charging infrastructure.[3,13,14] The EV contains a power cost that is "75% lower" than other forms of energy.[3] This leads to the operative costs of the EV being lower than those of the ICEV.[9] This is a major benefit of the EV, as the cost of fuel is a concern raised by the automotive market when considering the type of car to purchase. The only problem that arises is the location of charging stations.

The availability of charging infrastructure is a concern that most consumers have.[2,9,12] Buyers have the additional responsibility of locating the nearest charging stations, which, depending on the region, are not as widespread as gasoline stations.[9] In addition, EVs take a longer amount of time to charge,[2] leading to further inconvenience. However, some previous literature counters this claim on the direct correlation of charging infrastructure and EV uptake, stating that they are unrelated.[9]

The availability of charging infrastructure is worsened by the low range of the EV.[4] This correlates with the limitation of charging infrastructure, as the lower range results in increased charging frequency, leading to "range anxiety"[12] and excessive planning when traveling long distances.[9] Surveys conducted in a study covering the challenges faced by EVs claimed that the public saw mileage as the main limitation.[7] The survey proves that this limitation is a determining factor in the decision to purchase a car. The inconveniences caused by these limitations lead to a lower adoption of EVs in the industry.

Causes of the limitations

Production Costs and an increase in demand:

The mass production of EVs is aimed at reducing the emissions caused by the ICEV industry.[7,12] As the EV is powered by an electric powertrain, the basic components necessary for production are vastly different from those of a conventional vehicle. The mineral composition in the EV is "six times that of an ICEV",[9] demonstrating the complexity of the EV system in terms of raw materials in comparison to the ICEV. The relatively new technology utilized in the EV leads to the cost of production being vastly different from that of the ICEV.

The components of concern are within the powertrain, specifically the electric motor and battery. The cost of manufacturing a battery takes a considerable amount of investment relative to other components in the EV.[3,7,9] Most high-performance batteries and motors are manufactured using materials known as "rare-earth metals"[7] such as cobalt and nickel.[3] These metals have the benefits of increased battery performance, maximum output energy, increased range, and longevity of the vehicle.[7,9] The utilization of rare-earth minerals leads to the cost of raw materials taking up "50% to 70% of the total battery cost,"[9] which originates from the rarity and low yield of the metals.[7]

The large difference in supply in comparison to the demand leads to limitations in production costs. The demand for batteries is projected to "increase 6-fold from 2022 to 2030,"[15] resulting in increased demand for limited and low-yield metals.[7,15] This leads to a high cost of production, as the supply of materials used in the creation of high-performance EVs is outpaced by the demand of the market.

Supply Chain Disruptions:

The supply chain of EVs plays a pivotal role in the cost of production of the vehicle. The rise in demand for EVs leads to increased pressure on the supply chain.[15] When paired with a low supply,[7] the demand overrides the supply chain, leading to major shortages of materials.[9] This is worsened by the centralization of rare-earth mineral processing in one area,[9] such as the Tesla Mega-factory, resulting in delays due to the slow pace of sourcing the ores. This further causes a lower supply of materials, leading to a higher price of production.

The COVID-19 pandemic further exacerbated disruptions in the supply chain. This was demonstrated by the industry's reliance on semiconductors.[9] During the pandemic, the short-

age of semiconductors, due to underestimation of the demand,[9] led to losses of "7.7 million vehicle units in lost production and $210 billion in lost revenue in 2021 alone."[15] These major losses shed light on the lack of a secure supply chain in the EV industry, leading manufacturers to opt for cheaper, lower-quality materials.

One of the reasons suggested by Alshahapy et al. is the integration of the EV supply chain into the JIT (just in time) manufacturing method utilized by conventional car producers.[9] This method relies heavily on the overflow of a constant supply that the EV supply chain lacks.[9] This causes further complications in the EV supply chain, leading to the high upfront price of the car.

Secondary Data Analysis

Methodology Analysis:

Most papers that were analyzed in this study utilized one of two methodologies: the life cycle cost analysis (LCCA) and the total cost of ownership (TCO).

The life cycle cost analysis (LCCA) methodology is used to analyze the EV. This methodology falls under the category of Life Cycle Analysis (LCA), which covers 4 main areas that include goal and scope definition, inventory analysis, impact assessment, and analysis.[16] According to an LCA conducted by Raustad, the LCCA parameters were purchase price, fuel expenses, insurance, and registration costs,[17] demonstrating an economic focus of the methodology. Alternatively, Gao and Winfield's perspective on the defining parameters of the LCA contained the addition of GHG emissions and energy consumption, moving out of the scope of economic limitations or effects. Asiedu and Gu define the LCCA parameters as "costs to the manufacturer, user, and the society,"[18] generalizing their definition to include manufacturing and design costs of a product, widening the gap of the economic factors of the life cycle cost analysis, which they claim contradict each other, as the optimization of price reduces the aesthetic of the product and vice versa.[18] This demonstrates that the LCCA analysis methodology expands beyond the economic factors, which leads to a weakness as the other factors contrast with the economic parameters utilized to analyze products. The paper by Gao and Winfield further claims that "comprehensive LCA, specifically for advanced vehicles with commercial availability in today's market, has not been fully investigated,"[4] further illustrating the generalized approach towards parameters the LCCA utilizes, leading to insufficient data analysis on the cost of operations of the vehicle. This leads to further weakness in the methodology. However, contradictory to the claims of Gu et al., and Gao et al., Verma et al. define it as "The sum of total cost experienced during its lifespan,"[16] including parameters such as the cost of ownership and the cost of operations,16 which covers most of the economic implications of a product under the LCA methodology, giving a similar definition to Raustad. The different views on the methodology focus points of the LCCA give a sense of unreliability in the method of analysis on the economic front.

The Total Cost of Ownership (TCO) considers several parameters, depending on the authors' analysis and the specific areas the paper wants to discuss.[19] According to Gao and Windfield,[4] the TCO takes into account several parameters, including the initial cost of purchase and operational costs. The report by Alshahapy et al. expands the list of parameters to include "operation, maintenance, taxes and fees, financing, interest rates on loans, insurance, fuel, electricity, upgrades, and disposal."[9] These parameters cover the expenses of any type of conventional or electric vehicle. Sutakkul et al. describe the TCO's main objective as determining the current and future costs of a product, taking into account any assumptions.[19] This definition elaborates the purely economic scope of the TCO. Another format is utilized by Liu et al., which not only focuses on the economic implications of the vehicle but also on the social and environmental aspects. The TCO covers the economic implications of a product; however, depending on the assumptions chosen and other parameters taken into consideration, the TCO results may vary.

The TCO analysis covers the costs of operation in more depth than a life cycle cost analysis, and focuses on the economic parameters rather than a general analysis; this is the main methodology used to effectively analyze the parameters and was the main methodology used by most studies reviewed in this paper.

Analysis:

This paper utilizes secondary data analysis to further elaborate on the economic constraints of the EV through the range and find the correlation of EV sales compared to the economic strength of the region. As stated in the previous sections, one of the reasons that the EV is seen as inferior to the ICEV is the range, in contrast to the price, leading to low purchase rates. Furthermore, the high upfront price of the EV might serve as a barrier to citizens in lower-income countries. This additional information would be used to illustrate, firstly, the range in comparison to the price of the EV, the resultant effect on the purchase rates of the EV from a scope of economic strength, and the rate of ICEV sales in comparison to EV sales, dependent on the region. However, this data is highly exploratory and does not cover a large sample group.

Correlation between the price of EV and range:

The first data set analyzed was the purchase cost of several EVs in comparison to their range. My hypothesis was a positive correlation between the two variables. The higher the price of the EV, the higher the range of the EV.

Justifications:

This data analysis originated from the contrast between EVs and ICEVs in terms of range. The correlation between the price of the EV and the range of the EV demonstrates the value for money that the EV offers, and tests whether the price of the EV would shift the market to a position where the EV is comparable to the ICEV he importance of this study is to show how the association between range and price is seen and whether the EV fails to reach the market's expectation. The range per unit dollar comparison for each vehicle questions whether its battery capacity, before the inconvenience

DOI: 10.36838/IJHSR820.13

of charging, is worth the price. One of the main limitations that leads to low purchase rates is the low range and increased frequency of charging. Demonstrating the value of range for every dollar would show how the EV either contradicts or supports the issues raised by the market.

Method:

The main methods of analysis used in this first study were descriptive statistics of the EVs' range per unit dollar and a linear graph analysis with a confidence interval of the range compared with the prices of the EVs in the sample group. The analysis was done by researching the best-selling EVs of the years 2025, 2021, and 2018. Eight cars in each list were selected. The vetting of the EVs used in the study was based on the popularity rating of the vehicle from data blogs and the availability of relevant information on the vehicle from credible sources. Their price and range were researched from company data and car overviews online, and the data gathered was summarized in Table 1.1. Figures 1.1 and 1.2 were formulated by comparing the range and price of each model. The range per unit dollar was calculated by dividing the range by the price of each EV, and the bars were arranged in order of increasing range per unit dollar. Figure 1.1 is a linear graph of the line of best fit of the range as a function of the price of each of the vehicles, containing a shaded area representing the true mean values. Figure 1.2 is a visualization of each model's range per unit dollar.

Limitations:

Table 1.1.: Data set collected for each EV, including price, model year, and range to 5 significant figures. The data in the table captures more expensive vehicles having lower ranges than cheaper models.[20–25,25–42]

Vehicle	Year	Price/USD	Range/miles	Range per dollar/miles per USD
Tesla Model Y	2025	39990	321	0.00803
Chevrolet Equinox	2025	40790	319	0.00782
Geely Geomy Xingyuan	2025	9170	193	0.02105
Wuling Mini	2025	5570	75	0.01346
Chevrolet Bolt EV	2021	37495	259	0.00691
Volkswagen ID.4	2021	41190	190	0.00461
Audi e-tron	2021	66995	222	0.00331
Hyundai Kona	2021	49643	258	0.0052
Tesla Model S	2018	79200	259	0.00327
BJEV EC180/200 EV	2018	7146	96	0.01343
Nissan Leaf	2018	30885	150	0.00486
BYD e5 300/450	2018	32275	251	0.00778
Chevrolet Blazer EV	2025	46095	283	0.00614
Tesla Cybertruck	2025	80000	320	0.004
Honda Prologue	2025	48895	283	0.00579
Ford Mustang Mach-E	2025	39990	240	0.006
Hyundai Ioniq Electric	2021	34250	170	0.00496
Tesla Model X	2021	106440	335	0.00315
Porsche Taycan	2021	81250	199	0.00245
Tesla Model 3	2021	46440	272	0.00586
Kia Soul	2018	17095	111	0.00649
BWM i3	2018	48645	114	0.00234
Chevrolet Bolt	2018	37495	238	0.00635
Tesla Model S 100D	2018	97200	335	0.00345

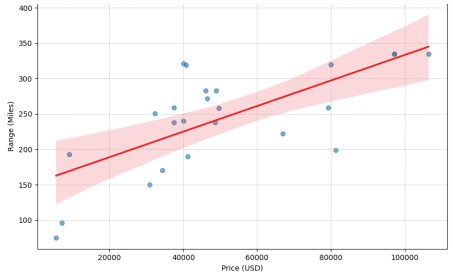


Figure 1.1.: A line graph of the price of the EV in comparison to the range of the EV using data from Table 1.1. The shaded area is the average mean or uncertainty surrounding the average value, and the confidence interval is 95%. The shaded area gets thinner as it approaches the mean region of the data. The blue dots are the actual values of the range and price of each vehicle, and the red line is the line of best fit. The graph shows a scattered correlation between the values, with many vehicle points deviating from the estimated mean.

The sample group of the vehicles analyzed was small due to several factors. Firstly, the data on the best-selling vehicles of each of those years was limited, with each list giving 10 vehicles that overlapped with other sources. Additionally, the data for some of the vehicles could not be found, leading to a low range of vehicles to choose from.

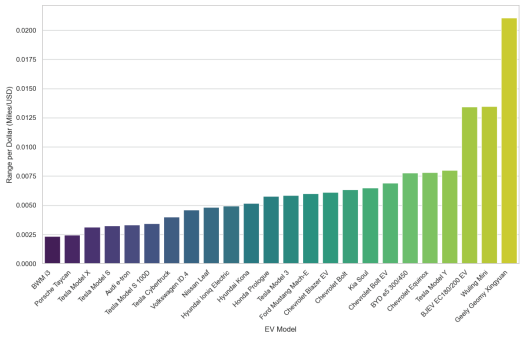


Figure 1.2.: A bar chart comparing the model of the EV and the range per dollar. All data is from Table 1. The graph shows cheaper models having higher values of range per unit dollar than expensive EVs.

EV sales in comparison to wealth distribution:

The second set of data analyzed was the EV sales in 14 countries: Brazil, the Netherlands, Germany, the United Kingdom, China, the United States, Canada, Chile, India, Israel, Japan, Mexico, Norway, and Italy. Their EV sales from 2015 to 2024 were compared to their GINI coefficient. Additionally, the ICEV sales in each country were recorded.

The GINI coefficient[43] is a commonly used parameter to evaluate wealth distribution. This parameter is utilized as a "measure of income inequality"[43] and contains a scale of 0-1. The higher the GINI index, the larger the income inequality in the nation. This was the parameter utilized to discuss the wealth distribution in relation to sales of EVs, to find a correlation between the two. I hypothesize that there is a direct correlation between EV sales and the GINI coefficient. Additionally, there would be a large difference between the sales of ICEVs and EVs.

Methods:

The main methods of analysis were a linear regression of the GINI coefficient and the EV sales of each country, and a visual representation of EV and ICEV sales demonstrated in a logarithmic graph. Data were sourced from Our World in Data on the EV sales, ICEV sales, and GINI coefficients of 14 countries. Table 2.1. summarizes this data. Figure 2.2. compares the wealth distribution and the EV sales, using a linear regression equation shown in Figure 2.1. to plot a dual bar graph. The GDP per capita and the urbanization rate of each country, sourced from the World Bank, served as the control variables used to formulate regression calculations seen in Figure 2.1. and Table 2.3. As another form of comparison, the sales of ICEVs and EVs in each country during the years 2015-2024 were sourced from Our World in Data, and the findings are summarized in Table 2.1. Figure 2.3. compares ICEV sales and EV sales in each country in a logarithmic graph. The countries were selected on an intended even distribution of economic

power and the availability of the data on EV and ICEV sales statistics.

Justification:

The study of EV sales in comparison to the GINI coefficient would show what association the EV sales have with each country's financial position. As the price of each EV affects the sale rate of the EV, the experimental data would show to what extent the EV sales are working towards the green effort if done on a larger scale. Additionally, the sales of the EV in comparison to the ICEV would show how low the EV sales are in comparison to the ICEV sales, proving the main thesis of this paper, that the EV has not reached large-scale adoption.

Regression analysis is important as the GINI coefficient is not the only known measure of wealth distribution and economic power of a country. Creating a graph that compares the EV sales vs the GINI coefficient of a country would be considering the assumption that the GINI coefficient is the only relevant determinant of economic strength and wealth distribution. The assumption would lead to a weakness in the analysis. Therefore, the use of linear regression using the GDP and urbanization rate as controls would allow for a more accurate analysis of the two factors.

The use of a logarithmic graph was to visibly show a distribution of EV and ICEV sales in a bar graph. Using the ICEV sales scale would lead to the EV sales bars becoming too small and the data not being easily visible. The logarithmic scale creates a balanced visual that accommodates the scales of the EV and ICEV sales.

Limitations:

Full Regression Formula and Symbol Definitions

Regression Equation:

EV_Sales = 6,190,199.80 + (33,538,733.10 × GINI) + (20.20 × GDP) + (-221,443.01 × Urbanisation)

Symbol Definitions:

- **EV_Sales:** Total Electric Vehicles sold (2015-2024)
- **GINI:** GINI Coefficient (Income Inequality index, 2021-2023)
- **GDP:** GDP per Unit Capita (USD, 2024)
- **Urbanisation:** Urbanisation rate (%)

Calculated Constants & Coefficients:

- Intercept (β_0): 6,190,199.7968
- GINI (β_1): 33,538,733.1034
- GDP (β_2): 20.1980
- Urbanisation (β_3): -221,443.0082

Figure 2.1.: The regression formula used to create the linear regression analysis demonstrated in Figure 2.2.

Table 2.1.: The data set collected from Our World in Data[44,45] that contains total EV sales in a decade compared with the GINI coefficient of each country (average from 2021-2023) and total ICEV sales in a decade. All the data on the sales are actual unit sales to 4 significant figures.

Country	Total Evs sold (2015-2024)	Average GINI coefficient (2021-2023)	Total ICEVs sold (2015-2024)
Brazil	215000	0.6918	18100000
China	33040000	0.5646	204300000
Germany	3470000	0.4712	30280000
Netherlands	819200	0.3888	3395000
United Kingdom	2127000	0.4738	21690000
United States	6095000	0.5869	157100000
Canada	806100	0.3111	13550000
Chile	11280	0.4304	2320000
India	238300	0.2551	19920000
Israel	225500	0.3789	3006000
Mexico	67460	0.4345	10890000
Japan	606400	0.3234	39770000
Norway	926400	0.2653	566200
Italy	603000	0.3434	15940000

The main limitation is the lack of many developing countries outside Europe, Southeast Asia, and North America in the study. Most nations have their GINI coefficient, GDP, and Urbanization rate documented on Our World in Data or the World Bank website. However, most developing countries do not have any data on EV and ICEV sales; therefore, the study only has a limited range of second- and first-world countries, with very few developing countries.

Table 2.2.: Data from the World Bank Organization,[46,47] including the urbanization rate in percentage and GDP per Capita of USD used as economic controls in the linear regression comparison model.

Country	Urbanisation rate/%	GDP Per Capita (2024)/USD
Brazil	88	10310.5
China	66	13303.1
Germany	82	56103.7
Netherlands	96	67520.4
United Kingdom	83	53246.4
United States	80	84534
Canada	83	54340.3
Chile	89	16709.9
India	35	2694.7
Israel	92	54176.7
Mexico	80	14158.8
Japan	92	32487.1
Norway	83	86785.4
Italy	70	40385.3

Table 2.3.: Linear Regression calculation table of EV sales as the dependent variable, GINI coefficient as the independent variable.

Variable	Coef.	Std.Err.	t	P>|t|	[0.025	0.975]
Intercept	6,190,199.7967	13,383,702.0916	0.4625	0.6536	-23,630,546.8171	36,010,946.4104
GINI Coefficient	33,538,733.1035	19,663,626.0431	1.7056	0.1189	-10,274,556.0541	77,352,022.2610
Urbanisation Rate (%)	-221,443.0082	178,665.0143	-1.2394	0.2435	-619,533.4681	176,647.4517
GDP Per Capita (USD)	20.1980	97.2238	0.2077	0.8396	-196.4301	236.8262

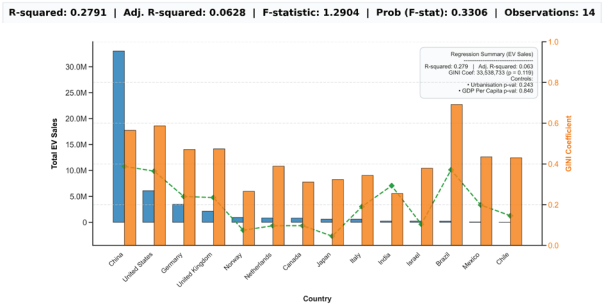


Figure 2.2.: A graph that compares the sales of EVs in each country and the nation's GINI Index; dashed lines are projections through the regression analysis; yellow bars represent the GINI coefficient, and the blue bars show the total EV sales. The results clearly show a mixed correlation of EV sales of each country vs their GINI coefficient.

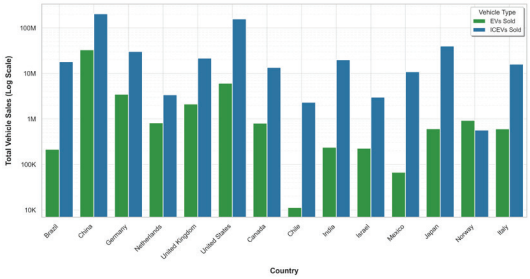


Figure 2.3.: A comparison of sales of ICEVs and EVs in the last 10 years (2015-2024) using a logarithmic scale. The graph shows a large gap between ICEV and EV sales.

Results:

Figure 1.1 is a linear regression line graph comparing the range and the price of EVs in the sample group. The line of best fit shows there is an estimated positive correlation between range and the price of the EV, and the red region

DOI: 10.36838/IJHSR820.13

indicates the true average mean values. 8 vehicles seem to have a ratio of range to price that falls under the average region, showing that a third of the vehicles analyzed follow the suggested trend. However, the number of extremes on the graph shows the association is not directly positive. 15 vehicle ratio values had a large range above or below the line of best fit. This indicates that the relationship displayed by the line of best fit is not proven to be positive. The data in the table supports the non-positive correlation, as it shows several vehicles collectively containing ranges that do not correlate with their value, such as the Volkswagen ID. 4 that costs 40,000 USD yet has a range of 190 miles, which is lower than the Geely Geomy Xingyuan at 193 miles with a price of 9170 USD. The extremes in the graph demonstrate an indirect correlation of the price as a function of range.

Figure 1.2. is a bar graph that compares each model's range per dollar ratio. The results show a mixed negative correlation between the price per unit dollar and the value of the model. More expensive models such as the BMW i3 ($48,645) and Porsche Taycan ($10,6440) have a low range per unit dollar (averaging at 0.002 miles per dollar), in comparison to the Geely Geomy Xingyuan ($9,170) and the Wuling mini ($5,570) having a higher range per unit dollar (0.021-0.013 miles per dollar). This shows that the range of the EV does not increase with price; instead, it decreases. However, this trend is contradicted by several vehicles such as the Tesla Model Y ($39,990) and Chevrolet Equinox ($40,790), which have a relatively high price in comparison to the cheaper models but are within the upper bound of range per unit dollar values (0.007-0.008 miles per unit dollar). This shows that the relationship between the range and price of each EV is not direct in this exploratory analysis.

Figure 2.2. is a dual bar graph containing a linear regression that compares the GINI coefficient and the sales of EVs in the period 2015-2024. The results vary depending on the nation; however, some of the sample countries seem to show a partial correlation between the GINI coefficient and EV sales. Countries such as the USA, Germany, and the UK have high values of the GINI coefficient and have relatively high values of sales of EVs, in decreasing order from the USA to Germany. This displays a positive trend in the GINI coefficient and EV sales, which is projected by the values shown by the regression line, and a negative trend in EV sales for countries with a lower GINI coefficient. However, seeing the trend in several countries contradicts the claim of a valid association. Nations such as Norway, the Netherlands, Canada, Japan, and Italy have variant GINI coefficients, with Norway visibly the lowest of the group and the Netherlands the highest. Nevertheless, all these nations have a similar number of EV sales, demonstrating that sales are not dependent on wealth distribution. Other countries, such as Israel, India, Brazil, Mexico, and Chile, have relatively high GINI coefficients in comparison to the rest of the sample group; however, the EV sales are the lowest in those countries compared to the rest of the group. This renders the association between the two factors very weak due to the extreme cases displayed by 11 of the 14 countries studied. China shows a great deviation from the correlation, as it has a GINI coefficient lower than the USA, but has more than triple the number of EV sales in the USA. This demonstrates a weak association between the two factors.

Figure 2.3. compares the sales of EVs and ICEVs in the years 2015-2024 in a dual bar graph with a logarithmic scale. The graph clearly shows the large margin between the EV and the ICEV sales. Most of the countries have ICEV sales above 10 million units, except the Netherlands, Chile, Israel, and Norway. This demonstrates a generally high ICEV purchase rate. The range of EV sales is all lower than 10 million units, except in China. This shows the ICEV sales over the last decade are extremely high compared to the EV sales. However, this is relative depending on the country. In several nations, such as the USA, China, Brazil, Germany, Mexico, India, Italy, Canada, Chile, the UK, and Japan, the purchase rates of ICEVs have a large margin compared to the purchase rates of EVs. In the Netherlands, Norway, and Israel, the rate of purchase of each car is almost equal. These differences in purchase rates show that the correlation in sales of both types of parameters is completely dependent on the nation in some cases; therefore, no direct connection can be determined from this exploratory data analysis.

Discussion

The first section of the Secondary Data Analysis compares the price as a function of the range of the EV. The summary derived from the data in both studies is that a weak and partial association exists between the two parameters. Several extreme values weakened the association between the range and price relationship, leading to the conclusion that range does not have a strong effect on the price of the EV. This trend could be caused by other factors affecting the price, such as the raw material costs and supply chain disruptors affecting the industry, as discussed in earlier areas of the paper. Additionally, different car manufacturers would focus on different needs when designing the EV. The results of the study are contrary to my hypothesis, as I predicted a direct and positive correlation between the two parameters. However, when taking into consideration Figure 2.2., there is a negative correlation displayed by the majority of vehicles studied. More affordable vehicles had a higher range per unit dollar in comparison to more expensive models, with a few outliers breaking the trend. This leads to the conclusion that the most expensive models do not exceed or equal the range in comparison to more affordable EVs in this study.

No papers were found that conducted such a comparison between EV range and price; therefore, no direct comparative analysis can be conducted. However, some papers made their own forms of analysis and inferences on the effect of the range on the price of an EV. The paper by Adeptu and Keshav comments on the limitation of range, stating that the solution needed to solve the limitation is to increase the size of the battery. The authors state, however, that the larger batteries are expensive, which would increase the cost of the car.[48] This shows an association between the two parameters, due to the mention that increasing the range would lead to an

increased production price. This claim contradicts the results found by this analysis, as more expensive models had a lower range. However, the expense of an EV is not directly tied to the use of advanced batteries that increase the range; therefore, the contradiction cannot be verified. Another study by Kim et al. states that both the range limitation and price of the EV are affected by the "high cost of battery."[13] This adds to the point made by Adeptu and Keshav, enforcing the perspective that both parameters are affected by a similar denominator, the size of the battery. This claim could be quantified if the sample group compared the cost of the battery used to the range rather than the upfront price of the EV. This perspective is countered by Woodley et al. The model used by Woodley et al in their empirical paper demonstrates that "a 1 percent increase in range is associated with a 0.423 percent decrease in price."[10] This elaborates a contrary perspective, that an increase in range leads to a decrease in price. This does align with the range per unit dollar study, as higher-priced cars had a lower range per unit dollar, showing that the statistics could support this claim if the association were tested over a larger sample group.

The second section of the analysis compares the GINI coefficient of a nation and the sales of EVs in the country over a decade (2015-2024). The results of the analysis show a weak positive correlation between the two parameters, which is greatly contradicted by several outliers in the sample group. An increase in the GINI coefficient led to an increase in EV sales in nations such as the UK, Germany, and the USA. Several outliers, such as China, the Netherlands, Italy, and Brazil, showed results that contradicted the trend, leading to the conclusion of a partially positive correlation. Other results were found in nations such as India and Norway, where both nations have similar GINI coefficients; however, EV sales in Norway were higher than in India. The hypothesis predicted that there would be a direct correlation. Several instances of outliers in the sample group counter this claim. The main result gained from the data is that wealth distribution is not a direct contributing factor to the sale of EVs, and other factors lead to the increase in sales. Such results could be explained by government intervention through EV incentive schemes, the availability of charging infrastructure, and external factors in each region. There are no papers that explore the wealth distribution as a function of sales of EVs; therefore, no comparative analysis can be formulated on these results.

The third set of data analyzed compares the sales of ICEVs and EVs in several countries. The results showed a minor gap in purchases over the decade, with nations such as the Netherlands, Israel, and Norway demonstrating minor differences between the sales of EVs and ICEVs. The majority of countries studied were outliers, such as the USA, Japan, Brazil, Germany, Mexico, India, Italy, Canada, Chile, the UK, and China, where the difference between the sales of ICEVs and EVs was significant. This leads to the conclusion that the ICEV is sold more than EVs, and the differences are dependent on the region. This trend could be explained by the economic efficiency in some nations due to the energy fuel mix used to fuel the cars, or enforced government policies that promote the EV industry. No other papers have created a graph plotting data on vehicle sales; therefore, no comparative analysis can be conducted. Yet, several papers discuss the gap between the ICEV and EV industries. Several authors, such as Bhat et al. and Wang et al., elaborate on the vast gaps between the sales of ICEVs and EVs. Both papers claim that the percentage of EVs in the car population as of 2019 was 1%,[7,47] and 16% in 2023.[12] This vividly demonstrates the gap between the EV industry and the ICEV industry, supporting the large differences in sales of the two vehicles in many of the countries in the sample group. Several other authors state the projection of EVs to increase by 22.5% from the year 2022-2030[9] or 5-fold from 2024-2029.[49] The increase in EV sales demonstrated by projections could explain the lower difference in the sales of both forms of vehicles in Norway, the Netherlands, and Israel. The projected growth shows an increase in the adoption of EVs.[14] Furthermore, a study by Costa et al. projects "parity between ICEV sales and EV sales will be achieved in 2037,"[14] which can be demonstrated in the example of the Netherlands, where EV sales in the last decade are nearly equal to the ICEV sales in the same period. The results of the analysis and readings studied show a steady increase in EV sales compared to ICEV sales, yet the outliers demonstrate that the EV industry still is not able to take over the ICEV market, despite the projections. However, a larger sample group of countries would be needed to verify this claim.

Solutions

Government Incentives:

One of the main limitations toward the large-scale adoption of the EV is the high purchase price of the vehicle.[9,10] There are several solutions suggested by other reviews to reduce the purchase price of the EV. Firstly, several governments have implemented fiscal incentives, such as tax reductions,[9,12] to increase the sales of the EV. Nations such as the USA have increased the federal tax credit income to boost the purchase rates of the EV.[9] Norway has seen significant progress through the use of federal incentives aimed at increasing consumer leniency towards the EV[12] leading to a 280% increase in purchase rates between the years 2018-2020.[48] Furthermore, India's investment of $1.2 million USD into the growth of the EV market through incentives and other projects,[47] shows further support for government investment effects on the rate of EV sales. Several studies indicate that Federal incentives have a direct positive correlation with the adoption of EVs by consumers,[12] with "the absence of studies" leading to the high price gap between EVs and ICEVs,[7] further limiting the sales of EVs.

This solution, however, is partially effective. The solution is completely dependent on the nation's enthusiasm towards EVs and its commitment to the electrification of the transportation sector.[12] "Some nations are skeptical"[12] about EVs, fearing a greener solution would damage their economy and threaten their industries. This lack of effort leads to further limitations in several regions,[12] lowering the rate of EV adoption. Additionally, when the "EV sales increase by a certain amount, the subsidies will be canceled."[7] This will lead to the price of the EV increasing, resulting in more limitations due to the mar-

DOI: 10.36838/IJHSR820.13

ginalization of the EV from the masses on the economic front. The gap in government intervention leads to EVs globally accounting for 1%[7] of the car population as of 2019.

Increased access to charging infrastructure:

The lack of charging infrastructure is another limitation of the EV that needs to be addressed. The low range of the EV leads to a more frequent need to fuel, inconveniencing buyers who are in regions with limited charging infrastructure.[4] Several studies show a direct correlation between an increase in charging infrastructure and the purchase rate of the EV, with Alshahapy et al. stating that the "investment in charging infrastructure is potentially more effective than financial incentives"[9] as a form of government intervention towards the sales of the EV, demonstrating the effectiveness of this solution. Some limitations with this solution include the location of the EV charging stations, permits and other legal parameters, and a proper source of electricity,[12] especially when aiming to have a larger EV fleet of vehicles in a country. However, some literature studies disagree with the direct correlation between lack of infrastructure and EV adoption.[9]

Supply chain disruptions:

As stated in the earlier sections of the paper, the supply chain disruptions caused by low supply of rare-earth materials, high demand, and concentrated processing lead to the increased price of EV production.[9] The supply chain disruptions are a major barrier to the adoption of EVs, as the initial cost of purchase is affected by the high cost of production.[7,9] To lower the main costs of production, several methods can be utilized to reduce the supply chain disruptions.

Firstly, one method is the use of alternatives to replace rare-earth metals in the battery and the permanent magnet. Materials such as ferrite are cheaper in cost compared to rare-earth metals, which are in low demand.[7] However, such materials contain limitations such as lower energy capacity and coercivity, reducing the overall quality of the permanent magnets.[7] The use of these materials leads to low-performance EVs with reduced range, further implementing its low adoption rate. In addition, battery alternatives such as Solid-State Batteries, Silicon Anode Batteries, and Lithium-Sulfur batteries are being adopted due to increased battery capacity and energy density, which increases the performance of the vehicle.[12] However, these solutions are still in the stages of research and trials, and are not yet able to be produced at a large scale due to safety concerns.[12]

Secondly, scientists are working on the disassembly of EV batteries for reuse to lower rare-earth metal demands. Methods such as pyrometallurgical recovery, physical separation, and robotic disassembly are being assessed to both decrease the life cycle impact of the EV and the cost of production of the vehicle.[14,49] This solution is not without its challenges. EV battery disassembly contains safety risks, such as explosions and chemical hazards, due to the handling of hazardous materials that pose a threat to the workers who conduct this process.[49] Additionally, the lack of data and technology needed to deconstruct these batteries, and the high cost of operation of these solutions, lead to the deconstruction of batteries being an inconvenience rather than a solution that can be adopted on a large scale.[51]

Furthermore, the decentralization of battery production would help increase the supply of batteries, resulting in lower prices for the EVs.[9] International collaboration and innovation are required to decrease pressure on the supply chain. However, this would require the handling of several political and social minefields associated with the mining and ore industries.[9]

■ Conclusion

This review covers the limitations that constrain the adoption of EVs on a large scale. The EV is yet to replace the demand for the ICEV, leading to a hindrance on the road to sustainability. The main limitations discussed were the high purchase price, marginalizing major consumers in the automotive industry;[7,9,12] lack of charging infrastructure, and lower range of the EV,[13,14] leading to customer inconvenience and further lack of interest from the market.

The reasons behind these limitations include high production costs of the battery, which takes up a significant part of the manufacturing costs.[3,7,9] Through the mass-scale use of rare-earth metals that are in low supply, the raw material costs take up to 70% of the expenses of battery production.[7,9] This is worsened by the centralization of processing the rare-earth materials, increasing supply chain disruptions.[9] This prompts the high purchase price of the EV, or manufacturers opting for cheaper alternatives, leading to EVs with lower ranges.[7]

The secondary data analysis conducted was highly exploratory and limited to covering comparisons of range and EV price, rate of ICEV and EV sales, and EV sales in comparison to wealth distribution. My hypotheses were a direct correlation between range and EV price, a direct correlation between wealth distribution, and a large margin between EV and ICEV sales within the nations analyzed. The results proved weak partial associations between the two factors (range vs price) and (wealth distribution vs EV sales), with outliers demonstrating the indirect effect of each of the compared parameters on sales and price of EVs. Additionally, the large gap between ICEV and EV sales differed regionally, proving wealth distribution is indirectly correlated with vehicle purchase rates. The secondary data led to the conclusion that several factors combined affect EV sales and prices. However, for more valid results, the studies need to be conducted with larger sample groups and more control variables taken into consideration.

The solutions suggested include government incentives, which have proven to be successful to some extent, with nations such as Norway reaping the benefits.[50] However, as stated in the paper, EV sales pertain to a minuscule part of the automotive industry, raising further questions of the profitability of increased incentives and charging infrastructure towards EV adoption. Additionally, the use of disassembly methods and rare-earth metal alternatives requires further research and investment to be adopted on a large scale.[12,48] Therefore, literature on potential alternatives to rare-earth metals is required to reduce production and price costs.

The EV is a solution that aims for a greener future. Several gaps in the field need to be addressed to improve the rate of EV purchase. One of the main gaps is the creation of effective government incentives and programs that propel the EV industry, through the study of countries that have such measures and their success rates. Additionally, the development of environmentally friendly and cost-effective batteries that increase the range of EVs needs to be researched to ease the manufacturing process and increase the market's confidence in EVs. The suggestions given would aid in advancing the EV field.

Acknowledgments

This paper was written under the guidance of Dr. Andrea Giusti and Mr. Marcin Kedziera through the IRIS program for research.

References

1. Ayodele, B. V.; Mustapa, S. I. Life Cycle Cost Assessment of Electric Vehicles: A Review and Bibliometric Analysis. *Sustainability* **2020**, *12* (6), 2387. https://doi.org/10.3390/su12062387.
2. V, H. S.; S, P.; J, A.; Dhivya, J.; Iyar, V. Data Analysis on the Comparison of IC vs EV Vehicles. In *2025, the 3rd International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA)*; IEEE: Coimbatore, India, 2025; pp 1–4. https://doi.org/10.1109/ICAECA63854.2025.11012171.
3. Kumar, A. A Comprehensive Review of an Electric Vehicle Based on the Existing Technologies and Challenges. *Energy Storage* **2024**, *6* (5), e70000. https://doi.org/10.1002/est2.70000.
4. Gao, L.; Winfield, Z. C. Life Cycle Assessment of Environmental and Economic Impacts of Advanced Vehicles. *Energies* **2012**, *5* (3), 605–620. https://doi.org/10.3390/en5030605.
5. Costa, C. M.; Barbosa, J. C.; Castro, H.; Gonçalves, R.; Lanceros-Méndez, S. Electric Vehicles: To What Extent Are Environmentally Friendly and Cost-Effective? – Comparative Study by European Countries. *Renew. Sustain. Energy Rev*. **2021**, *151*, 111548. https://doi.org/10.1016/j.rser.2021.111548.
6. Pero, F. D.; Delogu, M.; Pierini, M. Life Cycle Assessment in the Automotive Sector: A Comparative Case Study of Internal Combustion Engine (ICE) and Electric Car. *Procedia Struct. Integr.* **2018**, *12*, 521–537. https://doi.org/10.1016/j.prostr.2018.11.066.
7. Wang, Z.; Ching, T. W.; Huang, S.; Wang, H.; Xu, T. Challenges Faced by Electric Vehicle Motors and Their Solutions. *IEEE Access* **2021**, *9*, 5228–5249. https://doi.org/10.1109/ACCESS.2020.3045716.
8. Liu, Z.; Song, J.; Kubal, J.; Susarla, N.; Knehr, K. W.; Islam, E.; Nelson, P.; Ahmed, S. Comparing Total Cost of Ownership of Battery Electric Vehicles and Internal Combustion Engine Vehicles. *Energy Policy* **2021**, *158*, 112564. https://doi.org/10.1016/j.enpol.2021.112564.
9. Alshahapy, H.; Bozeman, J. F.; Carley, S.; Nock, D.; Matisoff, D. Modern Challenges Facing Electric Vehicle Adoption: A Review of Barriers to Adoption, Supply Chain Challenges, and Equity. *Environ. Res. Lett.* **2025**, *20* (2), 023002. https://doi.org/10.1088/1748-9326/ada6de.
10. Woodley, L.; See, C. Y.; Santos, V. R.; Yeo, M.; Palmer, D.; Nosenzo, S.; Nunes, A. Electric Vehicle Pricing and Battery Costs: A Misaligned Assumption. arXiv July 31, 2025. https://doi.org/10.48550/arXiv.2403.00458.
11. Nykvist, B.; Sprei, F.; Nilsson, M. Assessing the Progress toward Lower Priced Long Range Battery Electric Vehicles. *Energy Policy* **2019**, *124*, 144–155. https://doi.org/10.1016/j.enpol.2018.09.035.
12. Abdul Qadir, S.; Ahmad, F.; Mohsin A B Al-Wahedi, A.; Iqbal, A.; Ali, A. Navigating the Complex Realities of Electric Vehicle Adoption: A Comprehensive Study of Government Strategies, Policies, and Incentives. *Energy Strategy Rev.* **2024**, *53*, 101379. https://doi.org/10.1016/j.esr.2024.101379.
13. Kim, S.; Lee, J.; Lee, C. Does Driving Range of Electric Vehicles Influence Electric Vehicle Adoption? *Sustainability* **2017**, *9* (10). https://doi.org/10.3390/su9101783.
14. Costa, C. M.; Barbosa, J. C.; Castro, H.; Gonçalves, R.; Lanceros-Méndez, S. Electric Vehicles: To What Extent Are Environmentally Friendly and Cost-Effective? – Comparative Study by European Countries. *Renew. Sustain. Energy Rev*. **2021**, *151*, 111548. https://doi.org/10.1016/j.rser.2021.111548.
15. He, D.; Keith, D. R.; Kim, H. C.; De Kleine, R.; Anderson, J.; Doolan, M. Materials Challenges in the Electric Vehicle Transition. *Environ. Sci. Technol*. **2024**, *58* (28), 12297–12303. https://doi.org/10.1021/acs.est.3c09042.
16. Verma, S.; Dwivedi, G.; Verma, P. Life Cycle Assessment of Electric Vehicles in Comparison to Combustion Engine Vehicles: A Review. *Mater. Today Proc.* **2022**. https://doi.org/10.1016/j.matpr.2021.01.666.
17. Raustad, R. Electric Vehicle Life Cycle Cost Analysis. 2017 FSEC-CR-2053-17 https://rosap.ntl.bts.gov/view/dot/31875
18. Asiedu, Y.; Gu, P. Product Life Cycle Cost Analysis: State of the Art Review. *Int. J. Prod. Res*. **1998**. https://doi.org/10.1080/002075498193444.
19. Suttakul, P.; Wongsapai, W.; Fongsamootr, T.; Mona, Y.; Poolsawat, K. Total Cost of Ownership of Internal Combustion Engine and Electric Vehicles: A Real-World Comparison for the Case of Thailand. *Energy Rep.* **2022**, *8*, 545–553. https://doi.org/10.1016/j.egyr.2022.05.213.

Secondary Data analysis sources

20. 2020 Kona Electric Product Guide. https://www.hyundainews.com/assets/documents/original/42362-2021KonaElectricProductGuideSpecs080520pptx.pdf (accessed 2026-01-18).
21. *2025 Wuling Hongguang Mini EV 5 Doors New Models Larger Body.* https://wulingcars.com/product/2025-wuling-hongguang-mini-ev/ (accessed 2026-01-18).
22. *2021 Volkswagen ID.4 Review, Pricing, and Specs*. Car and Driver. https://www.caranddriver.com/volkswagen/id4-2021 (accessed 2026-01-18).
23. *2021 Chevrolet Bolt EV Review, Pricing, and Specs*. Car and Driver. https://www.caranddriver.com/chevrolet/bolt-ev-2021 (accessed 2026-01-18).
24. *2021 Audi e-tron / e-tron Sportback Review, Pricing, and Specs.* https://www.caranddriver.com/audi/e-tron-2021 (accessed 2026-01-18).
25. *2018 Tesla Model S Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/tesla/model-s-2018 (accessed 2026-01-18).
26. *2018 Nissan Leaf Review, Pricing, and Specs*. https://www.caranddriver.com/nissan/leaf-2018 (accessed 2026-01-18).
27. *The Best Electric Cars Of 2020 | Electrifying.* Electrifying.com. https://www.electrifying.com/blog/knowledge-hub/the-best-electric-cars-of-2020 (accessed 2026-01-18).
28. *The 10 Bestselling EVs of 2025 (So Far).* Car and Driver. https://www.caranddriver.com/news/g64540955/bestselling-evs-2025/ (accessed 2026-01-18).

DOI: 10.36838/IJHSR820.13

29. *Tested: 2018 Nissan Leaf.* https://www.caranddriver.com/reviews/a18660057/2018-nissan-leaf-test-review/ (accessed 2026-01-18).
30. *Model 3 – Sports Electric Sedan*. Tesla. https://www.tesla.com/model3 (accessed 2026-01-18).
31. *Geely Geome Xingyuan Latest News - CarNewsChina.com*. https://carnewschina.com/topics/geely-geome-xingyuan/ (accessed 2026-01-18).
32. *2025 Chevrolet Blazer EV Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/chevrolet/blazer-ev-2025 (accessed 2026-02-14).
33. *Cybertruck – Electric Utility Truck | Tesla.* https://www.tesla.com/cybertruck (accessed 2026-02-14).
34. *2025 Honda Prologue Review, Pricing, and Specs*. Car and Driver. https://www.caranddriver.com/honda/prologue-2025 (accessed 2026-02-14).
35. *2025 Ford Mustang Mach-E Review, Pricing, and Specs.* https://www.caranddriver.com/ford/mustang-mach-e-2025 (accessed 2026-02-14).
36. *2021 Hyundai Electric Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/hyundai/ioniq-electric (accessed 2026-02-14).
37. *2021 Tesla Model X Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/tesla/model-x-2021 (accessed 2026-02-14).
38. *2021 Tesla Model 3 Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/tesla/model-3-2021 (accessed 2026-02-14).
39. Author, T. *6 of the Best 2018 EVS, Both Full Battery Electric and Plug-In Hybrid | Sierra Club*. https://www.sierraclub.org/sierra/feature/6-best-2018-evs-both-full-battery-electric-and-plug-hybrid (accessed 2026-02-14).
40. *Best Electric Cars for 2018: Top-Rated EVs for 2018*. Edmunds. https://www.edmunds.com/electric-car/articles/best-electric-cars-2018/ (accessed 2026-02-14).
41. *2018 BMW i3 Review, Pricing, and Specs.* Car and Driver. https://www.caranddriver.com/bmw/i3-2018 (accessed 2026-02-14).
42. *2018 Chevrolet Bolt EV Review, Pricing, and Specs*. Car and Driver. https://www.caranddriver.com/chevrolet/bolt-ev-2018 (accessed 2026-02-14).
43. Hasell, J. Measuring Inequality: What Is the Gini Coefficient? *Our World Data* **2023**.
44. Hasell, J.; Rohenkohl, B.; Arriagada, P.; Ortiz-Ospina, E.; Roser, M. Economic Inequality. Our World Data **2023**.
45. *Number of new cars sold, by type*. Our World in Data. https://ourworldindata.org/grapher/car-sales?time=2013..latest&country=DEU~NLD~GBR~USA~BRA (accessed 2026-01-19).
46. *World Bank Open Data*. World Bank Open Data. https://data.worldbank.org/indicator/NY.GDP.PCAP.CD (accessed 2026-02-13).
47. *World Bank Open Data.* World Bank Open Data. https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS (accessed 2026-02-13).

Reviews

48. Adepetu, A.; Keshav, S. The Relative Importance of Price and Driving Range on Electric Vehicle Adoption: Los Angeles Case Study. *Transportation* **2017**, *44* (2), 353–373. https://doi.org/10.1007/s11116-015-9641-y.
49. Bhat, F. A.; Seth, Y.; Verma, A. Motivators and Barriers to the Widespread Adoption of Electric Four-Wheelers in India – A Discrete Choice Analysis of Potential Electric Four-Wheeler Buyers. *Travel Behav. Soc.* **2024**, *35*, 100748. https://doi.org/10.1016/j.tbs.2024.100748.
50. De Santis, M.; Silvestri, L.; Forcina, A. Promoting Electric Vehicle Demand in Europe: Design of Innovative Electricity Consumption Simulator and Subsidy Strategies Based on Well-to-Wheel Analysis. *Energy Convers. Manag.* **2022**, *270*, 116279. https://doi.org/10.1016/j.enconman.2022.116279.
51. Saxena, V. Enhancing EV Battery Lifecycle Management: Robotic Disassembly, Design for Disassembly, and Sustainable Solutions. *J. Energy Storage* **2025**, *119*, 116368. https://doi.org/10.1016/j.est.2025.116368.

Author

Kevi Sahara Ciku Karume is a student who aims to study Mechanical Engineering at a higher level and delve into the world of Automotive Engineering. She is passionate about integrating complex and high-performance systems into realistic and practical settings.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

A Comparative Laboratory Study of the Tensile Stiffness and Peel Strengths of Four Commercial Kinesiology Tapes

Lester Yew Suan Lim

Anglo-Chinese School (Independent), 121 Dover Road, 139650, Singapore; lesterlimyewsuan@gmail.com
Mentor: Rosalyn D. Abbott

ABSTRACT: Kinesiology tapes were selected for two functions in athletic care: providing mechanical support via the backing, and maintaining secure but removable attachment via adhesive peel strength. This study compared four tapes: Kinesio, KT, Decathlon, and an unbranded product, to identify which product best satisfied these factors. Tensile behavior was quantified by stress-strain testing of a section of tape specimen in the longitudinal direction, with linear regression of the elastic region to obtain effective Young's modulus. Adhesive performance was assessed using a 90° peel test on a silicon substrate with tape specimens under dry and wet conditions. Steady-state peel loads were recorded and normalized by specimen width. Bench testing used a calibrated force gauge and manual pull rate. The Kinesio tape had the highest Young's modulus, while the Unbranded tape had the lowest. The Unbranded tape had the highest dry peel strength, while the KT tape had the lowest. Wet peel strength decreased for every brand. Kinesio provides the highest backing stiffness and thus the greatest restoring force, being best for mechanical support. The Unbranded tape exhibits the strongest adhesion in both wet and dry conditions.

KEYWORDS: Biomedical Engineering, Rehabilitation Engineering, Young's Modulus, Peel Strength, Kinesiology Tape.

Introduction

In a study on university-level athletes, the prevalence of athletes with injuries was 91%, meaning that in a study with 81 athletes, only 7 athletes reported having no injuries. A total of 185 musculoskeletal injuries were reported, with an average of 2.28 injuries per athlete.[1] In athletes' recovery journeys for their sports-related injuries, they would use kinesiology tape for a range of reasons. In a study conducted with 200 students, 133 (67%) reported the usage of Kinesiology tape, with 121 (61%) of them applying the tape as a measure in dealing with musculoskeletal problems.[2] Kinesiology tape, or more commonly referred to as sports tape, is an elastic tape widely used in sports medicine and rehabilitation to provide adjunctive support for musculoskeletal injuries and to manage symptoms such as pain and swelling.[3]

These tapes are manufactured from an elastic fabric backing and a medical-grade, pressure-sensitive adhesive. The fabric of the tape can be stretched during application and subsequently would have a restorative force as it attempts to retract to its original length, which has been proposed to contribute to mechanical support.[4]

However, after trying a range of brands myself due to multiple ligament tears, I found out that there are limitations in all these brands, with each of them falling short in different departments. Some of the brands have an adhesive that is too sticky, to the point that it causes some tearing of the skin or red marks after removal. Some brands might be lacking in the bandage part of the tape, with it not having a strong enough tensile strength or not having a strong enough restoring factor for the tape to be effective.

Because commercial tapes differ in their fabric materials, weave, and adhesive formulation, a quantitative comparison that measures both the tensile behavior and peel adhesion is able to better inform consumers and athletes when selecting a tape for a given application. This study, therefore, compares four different brands of sports tape, Kinesio®, KT Tape®, Tarmak (Decathlon store brand), and an unbranded tape, using standardized mechanical metrics. More specifically, the tapes used are the Kinesio® Tex Classic (UPC: 850989002386), KT Tape® Original Cotton (UPC: 893169002813), Decathlon Tarmak Kinesiology Support Tape - Beige (Reference: 8596815), and the Unbranded tape (EAN: 5014749120696). The testing methods used will be a tensile restoring force test, used to derive the stress-strain behavior of the tape fabrics, and the 90° peel test on dry and wet substrates to mimic sweating, used to obtain the peel strength of adhesives.

However, due to the lack of professional testing equipment like a peel test machine, hand testing by pulling a force gauge introduces possible variability in displacement rate, alignment, and load application continuity. Additionally, due to ethical issues, the experiments will be done on a skin substitute, and while it may replicate most aspects and characteristics of human skin, it will not be a one-to-one replacement.

Literature Review

Kinesiology tape is frequently applied in clinical and athletic settings as an adjunct for pain control, proprioceptive facilitation, and functional support.[4] Systematic and randomized trials report mixed but sometimes favorable short-term effects on pain, range of motion, and perceived function in some pa-

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.14

tient groups, although the magnitude and consistency of these effects vary by condition and methodology.[5]

Kinesiology tapes are engineered for high elasticity, typically woven from cotton and elastic yarn, such as spandex,[6] which yields a soft, stretchy mesh that mimics the skin's tensile properties. For example, Kinesio® tape is advertised as 100% cotton with elastic fibers, allowing it to stretch up to 120% to 140% of its resting length, recoiling back to its original length when tension is released.[7] The material composition of the Tarmak tape is 97% cotton 3% elastane; that of the KT Tape® is 100% cotton, and lastly, the composition of the unbranded tape is 95% cotton 5% spandex.

The large stretch range enables the tape to exert a significant restoring force when stretched during application or movement, and can serve as an "external ligament", providing mechanical support, naturally going back to its original position to take the burden off an athlete's ligaments.[8] So, the elasticity of sports tape, quantified through its Young's modulus (E), the ratio of tensile stress to strain, would affect the mechanical support it provides.[9] The Young's modulus is the property of a material that dictates how easily it stretches and deforms, and is calculated through the equation $E=\sigma/\varepsilon$, where $\sigma=F/A$ and $\varepsilon=\Delta L / L_0$, where F is the force applied, A is the cross-sectional area of the tape, ΔL is the change in length, and L_0 is the initial length of the tape. Young's modulus is thus obtained from the gradient of a stress-strain graph and has the units Pa.[9] It is a physical quantity describing the ability of a solid material to resist deformation. The larger the value, the more force is required to deform the material.[10]

A tape with higher E will require more force to be elongated the same amount, therefore exhibiting a greater resisting force per unit elongation, increasing its capacity to resist joint translation and rotation when applied across a joint. The tape will thus be able to bear more external load and contribute greater resisting force across a taped joint, able to "share" load with injured ligaments or tissues rather than allowing that load to be taken entirely by the damaged structure.[11] So, taping can reduce injurious joint excursions and improve joint-repositioning accuracy, mechanisms by which stiffer external tape can help off-load injured ligaments and limit excessive motion.[11] Conversely, some tapes lack sufficient elasticity or tensile strength, so they exert less recoil and provide limited mechanical support.

Laboratory tests confirm that tape brands differ in tensile performance, with Kinesio® tape showing superior values in 3 out of 4 tensile metrics when compared to other tapes.[12] This thus proves how some brands of tape have higher tensile strength, being able to withstand greater load, as well as having a stronger restorative force and Young's modulus. So, different brands of tape would be expected to have varying levels of effectiveness in the support they provide. However, it is important to note that tape support is not static and that under repeated loading caused by factors like exercise, the tape loses elasticity. In fatigue tests of sports tapes, all specimens showed increased stiffness after cyclic loading. For example,[20] minutes of simulated use reduced support by 21% to 57% and made the tape up to 20% stiffer. So, these findings underscore that both the initial tensile properties and fatigue behavior of the tape material are relevant to how well it supports an athlete during prolonged activity. However, this study will only investigate the initial tensile properties of the different tapes.[13]

As for the adhesive of the tape, manufacturers coat the tape with a medical-grade acrylic adhesive, typically latex-free and hypoallergenic, so it sticks securely yet can be worn for several days.[6] For example, Kinesio® tape's adhesive is heat-activated and designed to stick firmly under tension, while remaining safe for sensitive skin.[6] This adhesive is usually applied on the tape in a patterned manner, allowing moisture to escape, permitting multi-day wear even during sweating. Manufacturers like Kinesio® claim that the tape will stay on through showering or exercise, simply drying quickly when wet.[7]

To test for the adhesive strength of the sports tapes, the peel test was used. The 90° peel test is a standardized method used to evaluate the peel resistance, or peel test, of an adhesive bond between a flexible adherent, like the sports tape, and a rigid substrate, where the flexible part is pulled perpendicular, or at a 90° angle, to the substrate surface.[14] The flexible tape is pulled upward at a constant speed, maintaining the 90° angle throughout, allowing for a steady-state plateau, in which the measured peel force becomes approximately constant with continued peeling.[15] This plateau is the region used to characterize the adhesive's practical peeling resistance. This setup enables measurement of the force needed to separate the bond under controlled conditions.[14] The peel force is then divided by the width of the bonded pad to obtain peel strength, which is the average force per unit width to separate a bonded adhesive,[16] to obtain a normalized property of the adhesive, which can then be compared to other adhesives.

The measured peel strength of sports tape adhesives varies by brand and condition. In standardized peel tests on dry and wet skin, the peak adhesive force has been reported to reduce adhesive grip by over half when applied on wet skin, as compared to dry skin.

Clinically, an adhesive on sports tape must strike a balance, resisting the movements and sweat of an athlete yet releasing without undue pain. If the peel strength is too high, removing the tape can cause skin injury. Medical adhesives with excessive adhesion are known to cause "skin tripping" and tears during removal, or even strip off layers of the skin or produce blisters, a known form of medical adhesive-related injury.[17] Conversely, if the adhesive is too weak, the tape will peel off prematurely and will not be secured when an athlete moves. Thus, the ideal tape adhesive provides moderate peel strength, strong enough to stay on through activity, but not so strong that it leaves the user with bruises or skin tears.

As research shows, there is a general trend of increasing pain with increasing peel force, and as such, the lower the peel force, the less pain experienced when the sports tape is removed.[18] The medical-adhesive review recommends a peel strength of 0.4–2 Ncm^{-1} for excellent removability, balancing secure attachment and painless removal.[19] As such, the ideal adhesive of a sports tape is to have a peel strength of 0.4–2 Ncm^{-1}, both on wet and dry surfaces, to ensure a painless removal while still keeping a secure attachment to the user's skin.

Methods

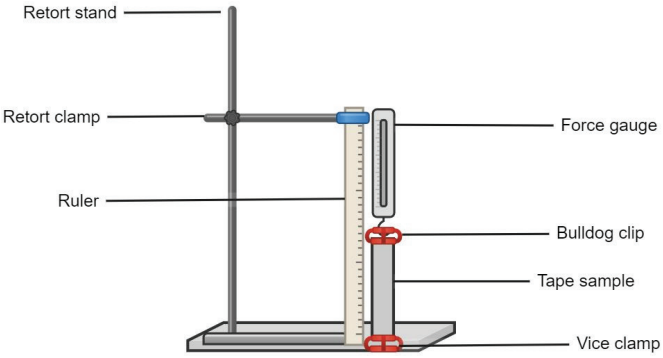


Figure 1: Set-up to test for the elasticity of tape samples (original image).

In this study, uniform strips of sports tape, 15.0 *cm* in length, and 5.0 *cm* in width and 0.05 *cm* in thickness, were prepared. These strips were kept in the same environment to ensure consistent adhesive behavior, minimizing variability and environmental impacts on their tensile behavior, ensuring that as many variables as possible are kept controlled to enhance accuracy. Each piece was then secured at one end using a vice clamp, firmly holding the tape in place, so that only a 10 *cm* section of the tape was exposed to testing (Figure 1). The opposite end was gripped by a bulldog clip attached to a calibrated force gauge. Next, a retort stand was placed beside the vice-clamp set-up, with a ruler being secured vertically next to the tape. This was to act as a vertical displacement scale, providing direct visual confirmation of the tape's vertical displacement and elongation when loaded and stretched. Thus, this set-up was used to measure the respective forces needed to elongate the tape by specific lengths, determining the tape's tensile characteristics.

Before each test, the force gauge was zeroed to ensure that there was no zero error, ensuring accurate readings. Tensile loading was applied to the tape by pulling up vertically on the force gauge to stretch the tape in the longitudinal direction to specified lengths, from 0 to 8 *cm* in 2 *cm* increments. The force reading was then recorded once stabilized, thereby capturing discrete data points that reflect the tape's elastic properties. Throughout testing, care was taken to maintain consistent clamp alignment, preventing slippage of the tape that could affect the accuracy of the stress-strain response.

The raw measurements of force F and elongation ΔL were then processed into engineering stress σ and strain ε, where $\sigma = F/A$ and $\varepsilon = \Delta L / L_0$, where A is the cross-sectional area of the tape and L_0 is the initial length of the tape. Then, σ was plotted against ε to obtain a stress-strain curve, which was used to determine Young's modulus.

To increase the precision of results and to account for the impact of random errors, 5 replicates were tested for each data point of each tape type to ensure statistical validity. Each replicate was carried out with a new piece of tape. The results were then reported as *mean ± standard deviation*, where error bars on the plotted graph represent the standard deviation of the data points.

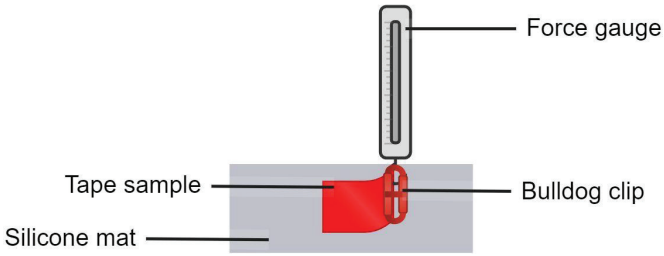


Figure 2: Set-up to test the peel strength of tape samples (original image).

To test for the peel strength of the adhesive on the different tapes, the tapes were cut into uniform rectangles, 5 *cm* in width and 7 *cm* in length. They were kept in the same environment to ensure consistent adhesive behavior, minimizing the discrepancy that temperature and humidity would cause between different tape samples.

A silicone mat was chosen as the test surface for the adhesives, due to its similar surface properties to human skin. The mat was cleaned with isopropyl alcohol and allowed to dry before testing, to ensure that there would be no surface contaminants that would bias results, improving repeatability and comparability between results.

A bulldog clip was attached to one end of the tape, so that the exposed area of the tape is 5 *cm* by 5 *cm*. This square area is then laid onto the silicone mat, smoothed with a single uniform pass of a flat-edged roller, and then left for a 10 *s* dwell time. A calibrated force gauge is then attached to the bulldog clip to measure the force required to peel the tapes (Figure 2). The specimen width was noted and later used to normalize the peel load to per-unit-width peel strength. Consistent application conditions, like pressure and dwell time, are essential as peel force is highly sensitive to initial contact conditions and effective bonded width. Thus, keeping the application conditions constant ensured that conditions were kept as similar as possible between samples, to reduce the impact of systematic error.

A 90° peel angle was established by using a set square that followed the edge of the tape as it was peeled. The set square was used as a calibration to ensure that the angle between the surface of the tape applied to the mat and the force gauge was being pulled to peel the tape off at a 90° angle. Before each test, the force gauge was zeroed and checked with a small known load to confirm calibration. Although laboratory-grade peel testers maintain exact kinematics, the manual arrangement described here follows the same geometry principles used in standardized 90° tests.

The peeling was also maintained at a controlled rate of 1.0 *mm/s*, and was maintained through the use of a ruler and a stopwatch, ensuring that 10 *mm* of tape was peeled. The force gauge then measured the constant force needed to peel the tape as the peel rate was kept constant. The peel rate was kept constant to ensure that the force required to peel the tape remains constant throughout the testing of one specimen. This is the steady-state peel region, which gives the most representative measure of adhesive peel strength, which was then normalized by specimen width to report peel strength.

DOI: 10.36838/IJHSR820.14

The adhesive strength was also tested on a wet surface, which aimed to investigate the adhesive properties under conditions of a sweaty skin surface. The silicone mat was dosed with 0.3 *ml* of deionized water, applied by spray from a distance of 5 *cm* away, to an area of 5 *cm* by 5 *cm*, immediately before the tape application. The same application conditions, such as pressure, dwell time, peel angle, peel rate, and data-logging procedure, were kept. The steady-state peel force was then recorded to investigate the effect of surface moisture, such as sweat, on adhesive performance.

Data processing was performed by normalizing the load by specimen width to obtain peel strength using the formula $p = F/w$, where p is the peel strength, F is the peel force, and w is the width of the tape sample. The peel strengths of the tapes were plotted against the brands to obtain a figure of comparison.

Five replicates were tested for each specimen under both dry and wet conditions, while maintaining all test variables constant, to enhance the precision of the results and minimize the impact of random errors. The results were reported as *mean* ± *standard deviation*, where error bars represent the standard deviation of the data points.

A one-way analysis of variance (ANOVA) test was then conducted separately for the peel strengths in both dry and wet conditions, in order to statistically compare the mechanical properties across the four tape brands. This was to confirm that the differences in mechanical properties were not due to random variation. Where significant differences were found, post hoc comparisons were performed using Tukey's HSD test to identify differences between specific brands. Statistical significance was set at $p<0.05$. However, the one-way ANOVA test was not possible for the Young's moduli of the tapes, due to a grouped data analysis being carried out to determine the Young's modulus of each tape. The average stress values of all five samples were found at specific strain points, and the average stress against average strain graph was plotted. Thus, a singular average Young's modulus value was found through the slope of this single, combined line. Since there was not one Young's modulus value for each replicate of each brand, but instead only one average value for each brand, the conditions required for an ANOVA test to be conducted were therefore not met.

Results and Discussion

Results:

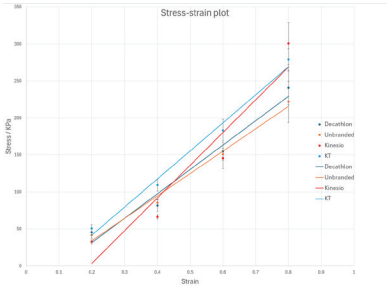


Figure 3: Linear regression of the stress-strain data, showing the elastic region of stress-strain curves. Kinesio tape had the steepest gradient, indicating Young's modulus was the greatest.

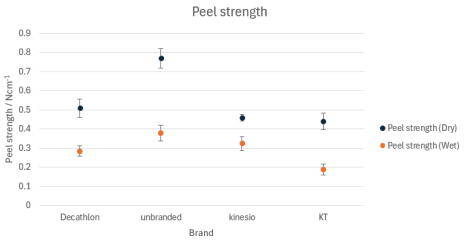


Figure 4: Peel the strengths of the respective brands of athletic tape. Dry peel strength was higher than wet, with the unbranded tape showing the highest values.

Table 1: Values of Young's Modulus. Kinesio's was the highest.

Brand	**Young's modulus** / KPa	R^2
Decathlon	329.59	0.9714
Unbranded	302.96	0.9893
Kinesio	442.19	0.9124
KT	379.46	0.9881

Table 2: Values of dry and wet peel strength. The unbranded tape had the highest values for both wet and dry conditions.

Brand	**dry peel strength** / Ncm^{-1}	**wet peel strength** / Ncm^{-1}
Decathlon	0.51 ± 0.05	0.28 ± 0.03
unbranded	0.77 ± 0.05	0.38 ± 0.04
kinesio	0.46 ± 0.02	0.32 ± 0.04
KT	0.44 ± 0.04	0.19 ± 0.03

The produced gradient of the linear regression of the stress-strain data is taken as the elastic moduli, or Young's modulus values, that differed across brands (Figure 3). These regression lines exhibited strong goodness-of-fit from the high R^2 values. In order from the highest to lowest Young's modulus values, Kinesio brand tape showed the highest modulus (442.19 kPa; R^2 = 0.9124), followed by KT (379.46 kPa; R^2 = 0.9881), then Decathlon, (329.59 kPa; R^2 = 0.9714), and lastly the unbranded product (302.96 kPa; R^2 = 0.9893) (Table 1). These fits explained more than 91% of the variance in measured stress over the tested strain range, thus supporting the use of the gradient of the linear regression models as a single effective modulus to compare backing stiffness between brands.

The peel testing also yielded distinct adhesion profiles in dry and wet conditions, and the values were plotted (Figure 4). Under dry conditions, the mean normalized peel strengths were, in order of highest to lowest, Unbranded (0.77 ± 0.05 Ncm^{-1}), Decathlon (0.51 ± 0.05 Ncm^{-1}), Kinesio (0.46 ± 0.02 Ncm^{-1}) and KT (0.44 ± 0.04 Ncm^{-1}). Through the one-way ANOVA test, it was determined that $F(3,16)=53.4$, and $p=1.48\times10^{-8}$, thus revealing a significant effect of tape brand on peel strength. Post hoc comparisons using Tukey's HSD test then showed that the unbranded tape had significantly higher peel strength than all other brands.

Wet substrate testing produced lower values for every brand, with the mean normalized peel strengths, from order of highest to lowest, being: Unbranded (0.38 ± 0.04 Ncm^{-1}), Kinesio

(0.32 ± 0.04 Ncm^{-1}), Decathlon (0.28 ± 0.03 Ncm^{-1}) and KT (0.19 ± 0.03 Ncm^{-1}). The percentage reductions in mean peel strength from dry to wet were, from highest to lowest, KT (57.4%), Unbranded (50.9%), Decathlon (44.5%) and Kinesio (29.8%) (Table 2). A one-way ANOVA test was carried out for the wet peel strength too, determining that $F(3,16)=21.9$, and $p=6.57\times10^{-6}$. Post hoc comparisons showed that the unbranded tape maintained a significantly higher peel strength than the KT and Decathlon tapes, while the differences in wet peel strength between the unbranded and Kinesio tapes were not statistically significant.

Discussion:

The tensile results indicate meaningful differences in backing stiffness between brands that are likely to influence a tape's mechanical support in use. Kinesio tape's substantially larger effective Young's modulus of 442.19 *kPa* implies that, for a given percentage elongation, it will generate greater tensile stress and therefore larger restoring forces than the other products tested. Conversely, the unbranded tape with a Young's modulus of 302.96 *kPa* is the most compliant and will produce the least passive resistance to deformation under equivalent stretch. According to literature, since restoring force underlies a tape's passive mechanical contribution to restricting excessive joint movement, these data suggest that Kinesio tape is the strongest candidate when maximal passive support is required, providing the most mechanical support and stability when applied, while the unbranded and Decathlon tapes will allow a greater range of motion under comparable application tension.

The noticeable inter-brand differences in their Young's modulus observed are most likely explained by the differences in textile architecture, like fabric weave, rather than the differences in constituent materials. The fabric backing of kinesiology tapes is textiles, in which factors like yarn crimp and weave geometry determine how load is transferred from macro-scale extension into fiber elongation, so two tapes made from similar materials can exhibit substantially different stiffness if their weave, yarn twist, or interlacing density differ.[20] In woven fabric mechanics, the weave structure substantially influences tensile stiffness. For example, plain weaves generally exhibit higher stiffness and strength due to denser interlacements and reduced porosity, while alternative weaves may alter the fabric's Young's modulus due to increased float and crimp levels.[21] As such, due to the relative material composition of each sports tape, the performance variation between them can be attributed to factors like fabric weave.

Next, adhesive performance showed a clear ranking and an apparent sensitivity to moisture. In the dry state, the unbranded tape had the highest mean peel strength, substantially exceeding the other three brands, with Kinesio and KT having similar dry peel strengths. The one-way ANOVA test conducted confirmed that the differences in peel strengths were statistically significant, with post hoc comparisons further showing that the unbranded tape had significantly greater peel strength than all other brands. Therefore, the unbranded tape demonstrated the strongest adhesive performance under dry conditions. Since the ideal adhesive of a sports tape is to have a peel strength of 0.4–2 Ncm^{-1} to ensure a painless removal while still keeping a secure attachment on the user's skin, and all brands tested have a peel strength below the upper limit of $2Ncm^{-1}$, this means that the tape with the best adhesive would be the Unbranded tape, with the strongest peel strength of $0.77Ncm^{-1}$.

However, exposure to a wet substrate reduced peel forces for every brand and altered the practical adhesion ranking in terms of in-use reliability. Although the unbranded tape retained the highest absolute wet peel load, its adhesion was reduced by 50.9%, and KT lost the largest percentage of adhesion, 57.4%. This leaves KT with the weakest wet adhesion with a wet peel strength of 0.19 Ncm^{-1}. Kinesio tape has the smallest reduction in peel strength from dry to wet of 30%, suggesting the greatest retention of adhesive function in moist conditions compared with the other tapes tested. In wet conditions, the unbranded tape still maintained the highest mean peel strength, with the one-way ANOVA test indicating that tape brand remained a significant factor affecting peel strength. Post hoc comparisons showed that statistically, the unbranded tape maintained significantly higher adhesion than KT and Decathlon, while differences between the unbranded and Kinesio were not statistically significant. Out of all 4 brands, only the Unbranded tape's wet peel strength lies within the range of 0.4–2 Ncm^{-1}that literature provided to ensure secure attachment, and as such, it is the best tape in terms of its adhesive strength, in both dry and wet environments.

■ Conclusion

The findings above have practical implications: clinicians and athletes should match tape selection to the intended function and conditions of use. If the priority is to provide maximum joint support, a high-modulus backing is preferable, and as such, Kinesio offers the best tape that provides the most mechanical support. If strong fixation in dry conditions is essential, a high-adhesion product like the

Unbranded tape is the best choice. Importantly, the substantial wet-state reductions observed here underscore the necessity of testing adhesives under conditions that mimic real use rather than relying on dry-state data alone. Thus, taking into account both the wet and dry peel strengths, the Unbranded tape has the strongest peel strength in both tests, and as such is the best tape to ensure a painless removal while still keeping a secure attachment on the user's skin. However, these recommendations should be interpreted in the context of the reported values and the study's laboratory conditions. Clinicians selecting tape should also weigh user comfort, skin sensitivity, and the specific demands of the sporting activity.

Limitations of this study include the modest sample sizes and the controlled bench protocol that cannot capture all *in vivo* variables like anatomical site, skin condition, motion patterns, and repeated application cycles. Additionally, peel and tensile behavior depend on test geometry and rate; thus, absolute values will vary with protocol. This is largely due to the limitations of manual testing. Using a handheld force gauge and manually pulling introduces variability in displacement rate, alignment, and load application continuity.

DOI: 10.36838/IJHSR820.14

As such, an improvement would be to use Universal testing machines, which, by contrast, offer precisely controlled crosshead speed, axial alignment, and feedback, which enhance accuracy and repeatability in modulus and peel measurements. Future work should expand, replicating numbers, include cyclic fatigue testing of backing elasticity, and perform in-vivo retention and skin-comfort assessments to link these bench metrics with functional and clinical outcomes. Additionally, while the inter-brand ranking by stiffness and adhesion in this study is informative and well-supported, the absolute modulus and peel strength values should be interpreted with caution and ideally re-validated using standardized, machine-controlled protocols.

AI Acknowledgments

During the preparation of this work, the author(s) used no AI tools.

Acknowledgments

I attest that the ideas, graphics, and writing in this paper are entirely my own. The author would also like to thank Rosalyn Abbott for her guidance and mentorship throughout the process of writing this research.

References

1. Lemoyne, J.; Poulin, C.; Richer, N.; Bussières, A. Analyzing Injuries among University-Level Athletes: Prevalence, Patterns and Risk Factors. https://pmc.ncbi.nlm.nih.gov/articles/PMC5596969/ (accessed 2026-04-24).
2. Krause, F.; Dust, K.; Banzer, W.; Vogt, L. Cohort Survey on Prevalence and Subjectively-Perceived Effects of Kinesiotape. *Deutsche Zeitschrift für Sportmedizin* **2017**, *2017* (05), 116–120. DOI:10.5960/dzsm.2017.272.
3. Cheatham, S. W.; Baker, R. T.; Abdenour, T. E. Kinesiology Tape: A Descriptive Survey of Healthcare Professionals in the United States. https://pmc.ncbi.nlm.nih.gov/articles/PMC8169012 (accessed 2025-07-15).
4. Boonkerd, C.; Limroongreungrat, W. Elastic Therapeutic Tape: Do They Have the Same Material Properties? https://pmc.ncbi.nlm.nih.gov/articles/PMC4868232 (accessed 2025-06-16).
5. Park, J.-S.; Yoon, T.; Lee, S.-H.; Hwang, N.-K.; Lee, J.-H.; Jung, Y.-J.; Lee, G. Immediate Effects of Kinesiology Tape on the Pain and Gait Function in Older Adults with Knee Osteoarthritis. https://pmc.ncbi.nlm.nih.gov/articles/PMC6855641/ (accessed 2025-07-16).
6. What Is Kinesio Tape? https://kinesiotaping.com/about/what-is-kinesio-tape/ (accessed 2025-07-16).
7. https://www.physio-pedia.com/Kinesio_Taping (accessed 2025-07-15).
8. Siegler, S.; Marchetto, P.; Murphy, D. J.; Gadikota, H. R. A Composite Athletic Tape with Hyperelastic Material Properties Improves and Maintains Ankle Support during Exercise. *Journal of Orthopedic & Sports Physical Therapy* **2011**, *41* (12), 961–968. DOI:10.2519/jospt. 2011.3476.
9. University of Birmingham. Physics A Level Revision Resource: Introduction to Young's Modulus. https://www.birmingham.ac.uk/study/undergraduate/schools-and-colleges/post-16/a-level-stem-resources/youngs-modulus (accessed 2025-07-16).
10. Long, H.; Lin, H.; Li, S.; Bai, Y.; Qin, L.; Xiao, T.; Qin, A. Nanomechanical Properties of CH4-Containing Coal during CO_2 Storage under Different Injection Pressures Based on Molecule Dynamics. *Applied Surface Science* **2022**, *590*, 153126. DOI:10.1016/j.apsusc.2022.153126.
11. Liu, Q.; Wang, L.; Dai, F.; Wang, G.; Chen, P. Effects of Ankle Kinesio Taping on Knee and Ankle Joint Biomechanics during Unanticipated Jumps in Collegiate Athletes. https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0305480& (accessed 2025-07-16).
12. Matheus, J. P.; Zille, R. R.; Gomide Matheus, L. B.; Lemos, T. V.; Carregaro, R. L.; Shimano, A. C. Comparison of the Mechanical Properties of Therapeutic Elastic Tapes Used in Sports and Clinical Practice. *Physical Therapy in Sport* **2017**, *24*, 74–78. DOI:10.1016/j.ptsp.2016.08.014.
13. BRAGG, R. W.; MACMAHON, J. M.; OVEROM, E. K.; YERBY, S. A.; MATHESON, G. O.; CARTER, D. R.; ANDRIACCHI, T. P. Failure and Fatigue Characteristics of Adhesive Athletic Tape. *Medicine & Science in Sports & Exercise* **2002**, *34* (3), 403–410. DOI:10.1097/00005768-200203000-00004.
14. Standard Test Method for 90 Degree Peel Resistance of Adhesives. https://store.astm.org/standards/d6862 (accessed 2025-07-16).
15. Vahdati, M.; Hourdet, D.; Creton, C. Soft Underwater Adhesives Based on Weak Molecular Interactions. *Progress in Polymer Science* **2023**, *139*, 101649. DOI:10.1016/j.progpolymsci.2023.101649.
16. Broughton, W. Testing the Mechanical, Therma,l and Chemical Properties of Adhesives for Marine Environments. *Adhesives in Marine Engineering* **2012**, 99–154. DOI:10.1533/9780857096159.2.99.
17. de Faria, M. F.; Ferreira, M. B.; dos Santos Felix, M. M.; Bessa, R. M.; Barbosa, M. H. Prevention of Medical Adhesive-Related Skin Injury during Patient Care: A Scoping Review. *International Journal of Nursing Studies Advances* **2022**, *4*, 100078. DOI:10.1016/j.ijnsa.2022.100078.
18. Chivers, R. A. Easy Removal of Pressure-Sensitive Adhesives for Skin Applications. *International Journal of Adhesion and Adhesives* **2001**, *21* (5), 381–388. DOI:10.1016/s0143-7496(01)00015-x.
19. Fitzgerald, D. M.; Colson, Y. L.; Grinstaff, M. W. Advancing Pressure-Sensitive Adhesives for Internal Wound Closure. *Nature Reviews Materials* **2022**, *8* (1), 3–5. DOI:10.1038/s41578-022-00516-y.
20. Hasan, K.; Luximon, A.; Luximon, Y. Comparisons of Kinesiology Tapes: Raw Materials, Fabric Structure, Physical Strength and Comfort Properties. *AHFE International* **2023**. DOI:10.54941/ahfe1003314.
21. Iftekhar, H.; Umair, M.; Hamdani, S. T.; Imran, S. M.; Nazir, M. S.; Ali, Z. Effect of Hybrid Weave Patterns on the Mechanical Performance of Woven Fabrics. *Journal of Natural Fibers* **2022**, *20* (1). DOI:10.1080/15440478.2022.2145411.

Author

Lester is a recent International Baccalaureate graduate from Anglo-Chinese School (Independent) who intends to major in Chemical Engineering. He combines hands-on experimentation with process-focused problem solving and is a recipient of multiple academic and research awards. He designed and ran experiments, analyzed data, prepared figures, and authored the manuscript.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Optimizing Ultrasonic Sensor Placement in Two-Sensor Systems for Reliable Navigation Using Dijkstra's Algorithm

Mahika Kasarabada[1], Devin Carroll[2]

1) The Lawrenceville School, Lawrenceville, NJ, USA; mahikakas2008@gmail.com
2) University of Pennsylvania, Philadelphia, PA, USA

ABSTRACT: Robotic systems limited to a few sensors focus on balancing sensor coverage with general reliability. In our two-sensor system, a critical area of focus is the angle between ultrasonic sensing cones, where narrow placement tends to limit peripheral awareness, while wider placement limits frontal awareness and creates frontal blind spots. Moderate placements around 60˚ are expected to achieve the best overall awareness. This study tests six inter-sensor angles (0˚, 30˚, 45˚, 60˚, 90˚, 120˚) on a simulated robot in ROS 2 paired with Gazebo Fortress that uses Dijkstra's algorithm to plan out a path. In this experiment, metrics including success rate (based on the number of collisions), path completion time, and path deviation were used to evaluate performance. Overall, moderate angles closer to 60˚ yielded the most promising outcomes, with paths that were more efficient and more in line with the optimal path laid out by Dijkstra's algorithm. Narrower angles tended to result in sideward impact, while wider angles led to delayed forward detection. Generally, these results demonstrate the significant impact that even small physical design choices can have on improving navigation in low-cost robotics– information that can inform the development of accessible and efficient robotics platforms.

KEYWORDS: Robotics, Autonomous Navigation, Ultrasonic Sensors, Sensor Placement, Path Planning.

Introduction

Navigation remains a critical challenge in autonomous robotics, requiring robots to respond to certain environments by accordingly planning routes and detecting obstacles nearby. Navigation is not only about getting from point A to B but also about ensuring that the robot can do so in a manner that is safe, efficient, and repeatable. Foundational tests confirm navigation as a basis for higher-level capabilities, including exploration, search-and-rescue missions, and even multi-robot coordination.[1] These higher-level functions are only made possible with reliable navigation.

To achieve robust navigation, platforms rely on systems with complex sensor arrays like LiDAR as well as vision, which offer rich perception but at a high cost. LiDAR works by emitting laser pulses and noting the time lapse between an outgoing pulse and the detection of the reflections to produce 3D maps.[2] Though these sensors enable high-accuracy localization and obstacle detection, sensors themselves can cost thousands of dollars, making them unattainable for hobbyist- or student-designed robots.3 On the other hand, vision-based systems provide insight into colors, textures, and features, though they require real-time processing for accurate image capture.4 Though these systems provide perceptual data, they often require more computational power in comparison to simpler robotics platforms. In practice, these limitations make advanced sensors suitable for grander projects like those conducted in research labs or in the industry, but unsuitable for classrooms, small-scale designs, and other cases.

Due to these constraints, ultrasonic sensors have emerged as a staple for low-cost and educational robotics purposes.[5] These sensors operate by emitting short bursts of sound above the standard human hearing range and then measuring the delay of the resulting echoes to estimate distance.[6] An ultrasonic sensor's conical detection field, as shown in Figure 1, allows for basic coverage of the frontal space and is easily mountable.[7] However, the simplicity associated with using these sensors comes with its own tradeoffs. Ultrasonic sensors are prone to environmental noise, limited effective range (not more than a couple of meters), and misreadings that occur when surfaces reflect sound away rather than back.[8] Thus, the placement of these sensors in a system becomes especially critical because poorly chosen angles can make the difference between a functional system and one that yields orientational issues, blind spots, and other features that degrade the robot's navigational performance.

The geometry of the placement of these sensors becomes critical. When two ultrasonic sensors are mounted in proximity to each other, their detection cones overlap, leaving the forward region accounted for, but the side areas neglected. This sensor configuration improves frontal detection, but degrades lateral detection, making the robot more vulnerable to side collisions. On the other extreme, when there is hardly any overlap between detection cones— a situation occurring with sensors placed far apart— a frontal blind spot is introduced.9 Both extremes are problematic, with narrow-angle placement sacrificing side coverage, and wide-angle placement sacrificing frontal awareness and delayed obstacle detection directly ahead. Studies emphasize placement geometry as a fundamental determinant of system reliability, as ultrasonic reflections and cone spread can amplify detection errors.[9]

This interplay is increasingly important when considered alongside path planning algorithms. Deterministic algorithms

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.15

like Dijkstra's are designed to guarantee optimal path planning.[10] If obstacles are left undetected due to poor sensor placement, then even the "optimal" path leaves the robot susceptible to collisions. Thus, without reliable sensing, even the most elaborate algorithms don't guarantee efficient navigation.

This issue is evident in educational robotics, in which students work with predesigned kits. Platforms like LEGO Mindstorms and VEX introduce ultrasonic sensors mounted in fixed positions.[12] While this eases beginner students into the design process, it introduces geometric limitations that students are unaware of contributing to the robot's performance. The result can be collisions or incorrectly linking this issue to "bad code" rather than suboptimal hardware placement. Identifying and addressing these limitations opens the door to an overlooked opportunity: with simple adjustments, designers can achieve better performance, reliability, and insight that can inform their takeaways and allow them to gain more educational value from low-cost systems.

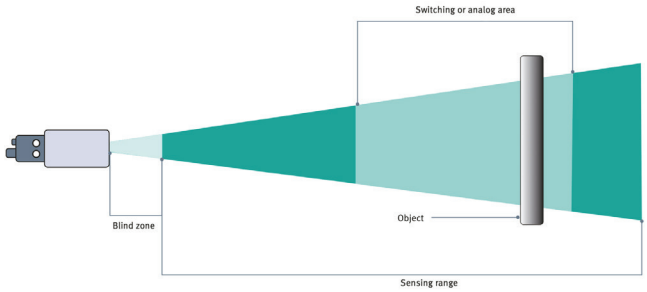


Figure 1: Ultrasonic sensor: conical field of view. This geometry presents the detection pattern of objects with respect to distance and angle.

While extensive research exists on advanced sensor and path planning, as well as a fair amount of research on sensor geometry, little prior work focuses on the effect of adjusting inter-sensor angles in minimalist systems, which in this case entails a simple box robot with two adjustable ultrasonic sensors. Most published work identifies sensor placement as a fixed design choice rather than an experimental factor. This gap is crucial for entry-level robots, as designs may be constrained to a minimal number of sensors and thus must optimize the placement of those sensors accordingly.

The objective of this study is to evaluate the effect of varied inter-sensor angles on the navigational performance of the robot in a controlled environment. A simulated robot with two ultrasonic sensors was tested across various trials in which angles were adjusted for a robot that maneuvered around a single obstacle and used Dijkstra's algorithm for path planning. Metrics, including success rate, path length deviation, collisions, and path completion time, helped quantify the robot's relative performance. We hypothesize that moderate angles between and near 45° and 60° will provide the most effective tradeoff between frontal and peripheral coverage.

Methods

This experiment was conducted in ROS 2 Humble simulation software paired with Gazebo Fortress. The robot, a simple box robot with two front-mounted ultrasonic sensors, was simulated on this software. Each sensor was simulated as a cone with a fixed field of view and range, and uniform sensitivity throughout. To simplify the experimentation process, the robot was not equipped with any actuators or wheels, and instead, path execution was computed on the grid that was discretized parallel to the direction of motion. This design choice helped to separate the effects of sensor placement from actuation variables.

Experimental Design:

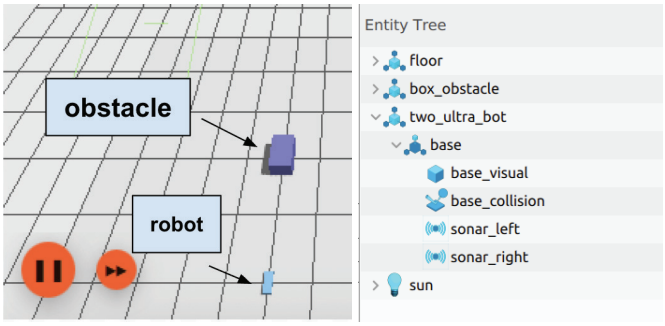


Figure 2: Simple ROS 2 simulation environment showing the simple robot, obstacle, and grid to traverse. (Left) The entity tree depicts a simple box robot with two ultrasonic sensors (not pictured in the simulation). By maintaining consistent environmental conditions across trials, performance differences can be attributed to shifts in sensor placement.

In this study, the independent variable was the inter-sensor angle (θ), which was tested at 0°, 30°, 60°, 90°, and 120°. The robot operated in a perpendicularly discretized grid environment as outlined in Figure 2, consisting of a cube-shaped obstacle, small enough for the sensors to detect it within their conical field of view, and a fixed starting and ending point. In this experiment, variables such as obstacle and robot size, grid resolution, and sensor range were all held constant across the trials. At each θ, six simulation runs were conducted. The reported results represent the average values of those runs to reflect consistent behavior, rather than outliers, across trials.

The navigation path was generated using Dijkstra's algorithm, which helped identify the shortest path to a goal. While the algorithm accounted for obstacles in the grid, it did not change based on missed sensor detections, thus ensuring that any differences in performance could be solely attributed to sensor placement instead of path computation variability.

Each inter-sensor angle was tested across various trials. The following metrics were recorded using ROS 2 logging tools to evaluate performance:

- **Success rate:** percentage of trials completed without collisions
- **Path length deviation:** difference between the actual traveled distance and the "optimal" path laid out by Dijkstra's algorithm
- **Collision count:** number of recorded collisions with obstacles
- **Completion time:** total time taken to reach the goal

Results

Across all runs, each sensor configuration registered at least one collision with the wall, indicating a uniform success rate of 0. Despite this, other observed metrics, including

path completion time and deviation, helped to distinguish between intersensor angle placement to provide meaningful insight (Table 1). For example, the lowest deviation (0.502) was achieved at an intersensor angle of 60°, and the fastest completion time (4.97s) was achieved at a sensor placement of 90°, indicating that optimal and time-effective paths were followed most closely at these moderate angles. In contrast, the 0° sensor configuration produced the longest completion time (93.27s) and the greatest deviation (1.126), confirming inefficiency when sensors are placed directly on top of one another. At the other extreme, the 120° sensor placement also proved inefficient, with equally high path deviation (1.126) and slower runs (6.00s). This is consistent with delayed frontal detection in wider placements.

Intermediate angle (between 30° and 90°) yielded mixed outcomes that generally proved to be more promising. At 30°, the robot followed an average path completion time of 6.78s with moderate deviation (0.711), while at 45° the robot took 8.61s, a greater amount of time, to complete the path, but at the same deviation (0.711). Quicker navigation did not correlate with the straightest paths. Such inconsistencies depicted that though moderate angles generally outperformed extreme ones, performance was variable and usually depended on what metric efficiency was measured by.

Table 1: Simulation data representing navigation performance metrics across varied intersensor angles, with each value representing the average across six simulation runs.

Success	Time (s)	Path Deviation	Collisions	Angle (degrees)
0	93.27	1.126	1	0
0	6.78	0.711	1	30
0	8.61	0.711	1	45
0	5.03	0.502	1	60
0	4.97	0.711	1	90
0	6	1.126	1	120

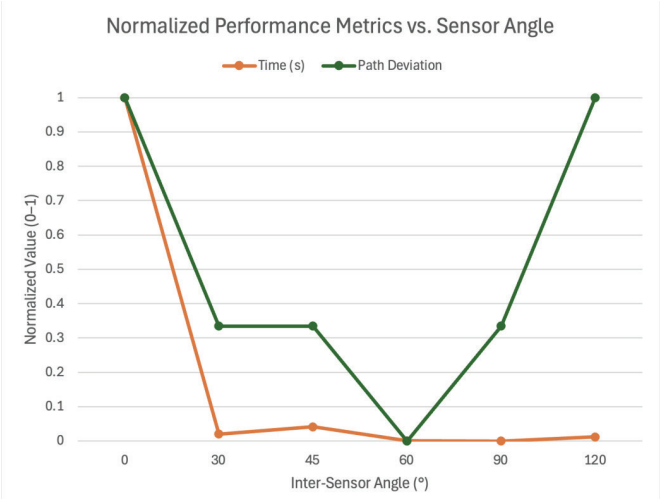


Figure 3: Comparison of normalized performance metrics— time and path deviation— against inter-sensor angle. Lower values indicate stronger performance, with 60° achieving the best overall balance across both metrics.

To properly enable cross-metric comparison, both path time and deviation were normalized on a 0-1 scale, with lower values indicating greater efficiency and performance (Figure 3). A sensor configuration of 60° consistently minimized deviation, and a sensor placement at 90° minimized completion time. Angles of 30° and 45° performed decently and occupied an intermediate range, while angles of 0° and 120° significantly underperformed across path time and deviation— the metrics that varied with each inter-sensor angle. Though an intersensor angle of 60° between ultrasonic sensors had the best tradeoff between performance metrics, the minimized completion time at 90° demonstrates that optimal placement can vary depending on the metric that is prioritized.

Discussion

The results support the hypothesis that moderate inter-sensor angles provide the most effective tradeoff between frontal and peripheral coverage. Though no trial fully avoided collisions, these efficiency metrics offered insight into the relative performance of the robot under various sensor configurations. At 90°, the robot achieved the lowest path deviation, and at 60°, the robot followed the most time-effective paths, indicating that these moderate placements allowed sensors to detect obstacles early enough to ensure a smoother and more effective trajectory.

By contrast, the 0° inter-sensor configuration suffered due to overlap: both sensors consistently monitored the area directly in front of them while neglecting the sides. This configuration created the illusion of redundancy but instead led to a gap in detection and thus long, inefficient runs. On the other extreme, an intersensor angle of 120° produced large blind spots directly ahead. The robot reacted late to the obstacles directly ahead of it, forcing abrupt adjustments to the robot's path and an inflated deviation. The geometry of beam overlap produces an explanation for this phenomenon: narrow angles sacrifice lateral vision, while wider angles prove unreliable in detecting frontal objects.

Even in this simplified simulation environment, such trends confirm that the placement alone of sensors (or ultrasonic ones, in this case) drove the navigational capacity and efficiency of the robot. These trends demonstrate that even with a deterministic algorithm such as Dijkstra's algorithm, which guarantees the optimal route, the robot's performance can be dictated by poor sensing geometry. Path planning alone cannot compensate for sensor misplacement.

In terms of low-cost robots in educational settings or hobby projects, this demonstrates how even small physical adjustments can have meaningful implications. In these budget-constrained settings, individuals may not have the ability to add to their system. With this understanding, individuals can optimize the robot's performance using the materials they already have. By intentionally selecting moderate inter-sensor angles, they can make minor practical adjustments that lead to greater results without any additional investment.

Conclusion

Overall, this study yielded the finding that adjusting the inter-sensor angle between ultrasonic sensors in a minimalist controlled environment influences navigational performance. This knowledge can be applied to two-sensor robotic systems

DOI: 10.36838/IJHSR820.15

and beyond. Moderate sensor placements, specifically at 60°, minimized path completion time and path deviation consistently, and vastly outperformed narrower and wider sensor placement configurations. Narrower configurations, such as at 0°, proved to be inefficient due to their disregard for lateral regions, while wider configurations, like at 120°, tended to create blind spots for the robot.

The implications of this study can be applied to designers of low-cost robotic systems, who may benefit from the understanding that geometric choices can shape system outcomes. By adjusting sensor placement, a minor but likely overlooked adjustment, designers can make steps towards smoother navigation and more efficient performance.

Future work should apply this knowledge on a broader scale and extend these tests into multi-obstacle environments, test the resilience of systems in the case of sensor failure, and validate whether these moderate inter-sensor angles are truly a "sweet spot" in real-world scenarios. Such experiments can further solidify the guidelines established by this study and apply this understanding to broader educational and practical contexts.

Acknowledgments

I would like to thank Dr. Devin Carroll and Teaching Assistant Kieran Tait for their mentorship and support during this project.

References

1. Siegwart, R. *Introduction to Autonomous Mobile Robots*, 2nd ed.; Intelligent robotics and autonomous agents series; MIT Press: Cambridge, Massachusetts, 2011.
2. Gregersen, E. Lidar | Scientific Technique | Britannica. *Encyclopædia Britannica*; 2020.
3. Wang, Z.; Menenti, M. Challenges and Opportunities in Lidar Remote Sensing. *Front. Remote Sens*. 2021, *2*, 641723. https://doi.org/10.3389/frsen.2021.641723.
4. Shahria, M. T.; Sunny, M. S. H.; Zarif, M. I. I.; Ghommam, J.; Ahamed, S. I.; Rahman, M. H. A Comprehensive Review of Vision-Based Robotic Applications: Current State, Components, Approaches, Barriers, and Potential Solutions. *Robotics* 2022, *11* (6), 139. https://doi.org/10.3390/robotics11060139.
5. Shaik, A. K.; Mohammadi, A.; Malik, H. A Systematic Review of Sensor Vulnerabilities and Cyber-Physical Threats in Industrial Robotic Systems. *IET Cyber–Physical Systems: Theory & Applications* 2025, *10* (1), e70023. https://doi.org/10.1049/cps2.70023.
6. Wong, J. *How Do Ultrasonic Sensors Work?* Beijing Ultrasonic. https://www.bjultrasonic.com/how-do-ultrasonic-sensors-work/.
7. Torres, C. R.; Abe, J. M.; Lambert-Torres, G.; da Silva Filho, J. I.; Martins, H. G. Autonomous Mobile Robot Emmy III. In *New Advances in Intelligent Decision Technologies*; Nakamatsu, K., Phillips-Wren, G., Jain, L. C., Howlett, R. J., Eds.; Kacprzyk, J., Series Ed.; Studies in Computational Intelligence; Springer Berlin Heidelberg: Berlin, Heidelberg, 2009; Vol. 199, pp 317–327. https://doi.org/10.1007/978-3-642-00909-9_31.
8. Borenstein, J.; Koren, Y. Histogramic In-Motion Mapping for Mobile Robot Obstacle Avoidance. *IEEE Trans. Robot. Automat.* 1991, *7* (4), 535–539. https://doi.org/10.1109/70.86083.
9. Hoomod, D. H. K.; Al-Chalabi, S. M. M. Object Detection and Angles Effectiveness by Ultrasonic Sensors HC-SR04. 2015, *6* (6).
10. Alshammrei, S.; Boubaker, S.; Kolsi, L. Improved Dijkstra Algorithm for Mobile Robot Path Planning and Obstacle Avoidance. *Computers, Materials & Continua* 2022, *72* (3), 5939–5954. https://doi.org/10.32604/cmc.2022.028165.
11. Julius Fusic, S.; Ramkumar, P.; Hariharan, K. Path Planning of Robot Using Modified Dijkstra Algorithm. In the *2018 National Power Engineering Conference (NPEC)*. 2018; pp 1–5. https://doi.org/10.1109/NPEC.2018.8476787.
12. He, S.; Hsieh, S.-J. Multi-Sensors for Robot Teaming Using Raspberry Pi and VEX Robotics Construction Kit. In *2018 ASEE Annual Conference & Exposition Proceedings*; ASEE Conferences: Salt Lake City, Utah, 2018; p 30829. https://doi.org/10.18260/1-2--30829.
13. Object, object. An Examination of Echoic Flow-Based Autonomous Guidance Using the Lego Mindstorms NXT Robot.

Author

Mahika Kasarabada is a rising senior at The Lawrenceville School whose projects range from AI-driven wellness apps to pneumonia detection models. At school, she is involved in Girls Who Code, FIRST Robotics, and the AI Council. Beyond academics, she dedicates her time to singing, playing tennis, and serving her community.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

The Effect of the Reduction of the Cash Payment Limit on Compliance (VAT gap) in France

Maria Sultanova

International School of Nice, 15 Avenue Claude Debussy, Nice, 06200, France; msultanova777@gmail.com

ABSTRACT: The cash payment limits in Europe are getting increasingly stringent due to the lack of traceability. Cash payments often concern governments because of their use for money laundering, terrorism financing, and other illegal activities that increase countries' shadow economies. For instance, one of the noticeable policies, reducing the cash payment limit to €1,000, was established in 2015 in France. In this paper, we explore the extent to which this policy affected tax compliance in France by using the difference-in-difference method together with empirical data from relevant papers for comparing the VAT gap before and after 2015 in France and Germany. We expect to see a decrease in the VAT gap, reflecting that the cash payment limit was consistent with the reduction of informality in France. Our difference-in-difference estimate shows that the VAT gap as % of VTTL after 2015 declined in Germany by 2.8 pp more than in France. Although the VAT gap decreased in France after 2015, the decrease was smaller than in Germany. Therefore, the difference-in-difference estimate suggests that the policy had a limited relative effect on tax compliance. Consequently, these observations could be taken into consideration for future policy studies and decision-making for the EU economies.

KEYWORDS: Economics, Public Finance, VAT Gap, Tax Compliance, Cash Payment Limits.

Introduction

The European economy shifts a lot from what is often called "traditional" cash payments to more traceable, online transactions, combating informality. One of the primary reasons for this is the increasing VAT gap that many governments face. Low compliance suggests multiple issues, not only for the government but also for society as a whole. It decreases VAT revenues, reflecting the increasing informal sector of economies. Even though the shadow economy is difficult to measure or estimate, the VAT gap is widely considered to be one of its main determinants.[1] Evidently, following the VAT gap estimations,[2,3] some countries need government intervention due to the low levels of compliance. France, for instance, has one of the largest VAT gaps in the 2009-2013 time period in the EU. Moreover, the country experiences challenges in VAT tax collection. VAT receipts in France appeared to be only 14.5% of total tax revenue in 2014, while the EU average was 17.5%.[4] That is why this country serves as a prominent case study of the EU member state that has implemented one of the strict cash-payment limit reductions (see Appendix A: Table A1).[5] As the OECD notes, "In France, since September 2015, there has been a limit of €1000 for cash payments between businesses and individuals as well as business to business transactions."[1] This means that residents of France wishing to purchase by cash can pay a maximum amount of €1,000 in cash. Many consumers are highly dissatisfied with such requirements, as the system doesn't allow for even a slight exception, making purchasing a luxury good a more complex action today.

However, the cash-payment limitations are essential in the EU, because they often control terrorism finance and money laundering.[6] Cash payments are generally untraceable, which makes them a driver of illegal and informal activity.[1] Another piece of evidence from the paper on the determinants of compliance, indicating greater improvement, states that "the fiscal evasion risk, which includes that of VAT evasion, is larger in less digitalized countries."[7] This implies that traceability is crucial for keeping low levels of informality. Additionally, the same paper validates the strong relationship between VAT gap and shadow economy. That is why we link cash payment limitations with informality, assuming that these restrictions influence shadow economies directly, helping to reduce them in particular. Consequently, many European countries, including France, aim to control cash flow. Nevertheless, despite the decrease in demand, cash remains essential as well as highly accessible in the French economy.[8] The consumers' behavior is often well explained by the cash-holding model, which suggests that "the agent pays cash whenever he has enough cash and with another payment instrument otherwise."[9]

We investigate how the 2015 reduction of the cash payment limit to €1,000 affects tax compliance in France. In our research, the majority of the literature on the cash payment restrictions doesn't explore or focus on any economies in particular, but on the EU in general.[10] While previous studies have analyzed VAT gaps across the EU, fewer studies focus on specific country-level policy changes. One of the relevant examples would be an empirical paper that raises a similar question of whether the cash restriction affects the tax evasion levels, investigating the economy of Italy.[11] The Italian case findings show that cash restriction policies rather change the tax compliers' behavior than reduce tax evasion itself.[11] Our paper contributes to economic literature in three ways. First, we extend the literature by studying the impact of a specific cash payment limit reform in France, focusing on one country rather than the EU as a whole. Second, we use a difference-in-difference approach with Ger-

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.16

many as a control group, which has not been applied in this context in the existing literature. Third, we focus directly on the VAT gap as the outcome variable to measure tax compliance.

Being one of the largest economies in the EU, France appears not only a strong representative, but also an important one to explore in order to understand the link between the earlier-mentioned policy and the French economy. The strictness of cash payment limitation in France makes it illustrate such policies better. Furthermore, identifying the effect is crucial for policymakers' and governments' decisions against the shadow economy. We study the behavior of the VAT gap, which can potentially reflect a country's informal sector, since it has a direct and exceptionally strong correlation with the VAT gap.

We compare compliance in France before and after the 2015 cash-payments limit reduction. We expect compliance to increase (or, in other words, the VAT gap to decrease) as payments should become more traceable. This might leave less freedom for informality and more overall control. This trend is already observed in related economic literature, as the EU's VAT gap has declined in 2018-2020.[3] It also might suggest the significant contribution of France, as the time period acquired is after the restriction implementation, which satisfies our requirement for testing the hypothesis. However, it is important to highlight that we consider the cash-payments limitation to be one, but not the only factor that influences tax evasion. For example, changes in labor informality might also affect tax revenues and thus the VAT gap itself both positively and negatively.[12]

Our hypothesis is going to be tested with the use of recalculated empirical data from reports on the VAT gap in the EU-28 Member States and a strong review of already existing literature in the field.[2,3] The difference-in-difference method shall be performed on the two separate economies (France and Germany), respectively representing one treated and one control.

Our paper is then structured as follows. In the next Section (2), we present our data and methodology with resources. Next, in Section 3, the results are only shown and described, but are discussed later, as well as evaluated in Section 4. Finally, Section 5 will conclude the paper. Some of the sections are divided into subsections for smoother guidance and a breakdown of the paper.

Methods

Data:

In our paper, we used a variety of economic literature and empirical data from economic reports to test the hypothesis. In this subsection, we briefly list and explain the main sources that were used for collecting and evaluating data. We chose the VAT gap as our key variable because of its direct relationship with the shadow economy.[13]

We compared the data on the VAT gaps in the EU by taking the examples of the French and German economies in the "reports of VAT gaps in the EU-28 member states" of 2016 and 2020, respectively, which were our main data sources.[2,3] The first paper reflects a 2010-2014 time period, which is before the 2015 cash-payment limit implementation in France, while the second is a follow-up update with the data when the policy was already in place. Germany was chosen to be a control group for several reasons. First, no similar cash restriction policy took place over the observed time period in Germany. This satisfies our difference-in-difference approach. Second, both French and German economies are considered to be among the largest EU economies, making them "together account for almost half of the euro area's GDP."[14] With their technological advancements and world influence, the two economies appear to be comparable.[14] Third, Germany's compliance gap is close to the EU's average, which makes it a strong representative not only of a specific economy, but also of a more collective model of the EU economies.[15] So, we observe the same time period for the two developed economies in the EU to construct the correlation between changes in VAT compliance in France and its cash payment limit reduction.

In order to explain the correlation between VAT gap and cash payments, we used data on both the German and French economies. To show the relationship between the size of the shadow economy and VAT gap, we recorded shadow economy estimates as percentages of official GDP for France and Germany from the "Taxation of the Informal Economy in the EU" report.[16]

Methodology:

This study explores the effect of the cash payment limit reduction of 2015 in France on its compliance, through the difference-in-difference method. We use data from reports on VAT gaps in EU member states and focus on France and Germany over the period of 2010–2018.[2,3] The pre-policy period is 2010–2014, while the post-policy period is 2015–2018. This data allows us to compare VAT gap behavior before and after the policy. Over the same time period (2010-2018), we observed the behavior of the four key variables: VAT gap, VTTL, VAT revenues, and VAT Gap as a percentage of VTTL, which helped us to evaluate the effectiveness of the cash payment limit.

The main variable used in the analysis is the VAT gap as a percentage of VTTL. This measure allows for comparison across countries by removing the effect of differences in economic size. It is commonly used as an indicator of tax compliance. In order to eliminate the problem of data inconsistency that arose when using the data from two separate reports due to the revised methodology of the 2020 report, the data before 2014 were spliced using the ratio method (see Appendix A for raw series).[2] Our calculations harmonized the data, removing these inconsistencies and allowing for accurate comparison. We ran a regression to evaluate the statistical significance of our results (Table 4). We also computed the difference-and-difference estimate with a robustness check (Tables 3, 5). To evaluate the effect of the 2015 policy, we apply a difference-in-differences approach. France is treated as the policy country, while Germany is used as the control group. We compare the change in the VAT gap before and after 2015 in both countries to estimate the policy effect.

This paper explicitly highlights the link between the VAT gap and informality by calculating the correlation between the two variables. This allowed for a deeper analysis, representing not only the consistency of the policy with an improvement of the VAT gap, but also the significance of the 2015 cash payment limit. In the concluding section, we suggested potential improvements as well as stated the limitations of our method.

Difference-in-Difference Approach:

To apply the difference-in-difference method, we investigate the behavior of the VAT gaps in France as a treated group and Germany as a control group over two different time periods: the pre-policy period (2010-2014) and post-policy period (2015-2018).

The difference-in-difference estimate is computed as follows:

$$\delta_{DiD}^{2010-2018} = \left(\overline{Y}_{FR,2015-2018} - \overline{Y}_{FR,2010-2014}\right) - \left(\overline{Y}_{GR,2015-2018} - \overline{Y}_{GR,2010-2014}\right)$$

(Equation 1)

Where:,

$\delta_{DiD}^{2010-2018}$ - the DiD estimate value for VAT gap as a percentage of VTTL

$\overline{Y}_{FR,2010-2014}$ - the mean of VAT gap as a percentage of VTTL for the pre-policy period (2010-2014) in France

$\overline{Y}_{FR,2015-2018}$ - the mean of VAT gap as a percentage of VTTL for the post-policy period (2015-2018) in France

$\overline{Y}_{GR,2010-2014}$ - the mean of VAT gap as a percentage of VTTL for the pre-policy period (2010-2014) in Germany

$\overline{Y}_{GR,2015-2018}$ - the mean of VAT gap as a percentage of VTTL for the post-policy period (2015-2018) in Germany

A positive difference-in-difference estimate indicates that the reduction in the VAT gap in France was smaller than in Germany, while a negative value would indicate a stronger improvement in France.

We then run a regression in order to estimate the treatment effect, evaluating the statistical significance of our difference-in-difference estimate.

Our regression model is constructed as follows:

$$Y_{it} = \beta_0 + \beta_1 France_i + \beta_2 Post_t + \beta_3 (France_i \times Post_t) + \varepsilon_{it}$$

(Equation 2)

Where:

Y_{it} - the VAT gap as a percentage of VTTL

$France_i$ - Dummy variable that identifies the treated country (Germany is 0, France is 1)

$Post_t$ - Dummy variable that identifies the policy period (2010-2014 is 0, 2015-2018 is 1)

β_0 - Baseline VAT gap for Germany before 2015

β_1 - Difference between France and Germany for the pre-policy period (2010-2014)

β_2 - Average post-policy period change affecting both France and Germany

β_3 - the DiD estimate value for VAT gap as a percentage of VTTL

($France_i \times Post_t$) - An interaction term that is used to estimate the difference-in-difference treatment effect

ε_{it} - Error term capturing unobserved factors affecting the VAT gap

β_3 represents the difference-in-difference estimate. A positive value of β_3 would indicate a smaller reduction in VAT gap in France relative to Germany, while a negative β_3 value would indicate the opposite.

The regression model also provides standard errors, which will allow us to make a judgment on the statistical significance of the result. The model is estimated using the Ordinary Least Squares (OLS) regression framework.

In our paper, the difference-in-difference approach has an assumption that the economic behavior of the two countries would have been similar if no policy had taken place. In other words, it relies on the parallel trends assumption. Our visual presentation of the VAT gap as a percentage of VTTL in both countries shows that the pre-policy trends in the two countries are not perfectly parallel. (Figure 4) This weakens the validity of our difference-in-difference estimate. Because the initial pre-policy trends of the two economies differ, Germany does not fully satisfy the assumed model of the French economy without government intervention. Nevertheless, there is a noticeable trend of the VAT gap as a percentage of VTTL increasing in France and Germany before the policy implementation in 2015. Therefore, they should be interpreted cautiously, with the mentioned limitation taken into consideration.

In our robustness check, we applied the same treated and controlled groups, but with slightly different time periods: pre-policy period (2010-2014) and post-policy period (2016-2018), which excludes the year of the policy implementation. We use an alternative post-policy period for our robustness check in order to eliminate the possible effects, as well as the sensitivity of the transition year (2015). A significant change in the difference-in-difference estimate after the robustness check might suggest the dependency of the results on the timing.

The robustness check was therefore computed as follows:

$$\delta_{DiD}^{2010-2018} = \left(\overline{Y}_{FR,2016-2018} - \overline{Y}_{FR,2010-2014}\right) - \left(\overline{Y}_{GR,2016-2018} - \overline{Y}_{GR,2010-2014}\right)$$

(Equation 3)

Where:

$\delta_{DiD}^{2010-2018}$ - The DiD estimate value for VAT gap as a percentage of VTTL.

$\overline{Y}_{FR,2010-2014}$ - the mean of VAT gap as a percentage of VTTL for the pre-policy period (2010-2014) in France

$\overline{Y}_{FR,2016-2018}$ - the mean of VAT gap as a percentage of VTTL for the post-policy period (2015-2018) in France

$\overline{Y}_{GR,2010-2014}$ - the mean of VAT gap as a percentage of VTTL for the pre-policy period (2010-2014) in Germany

$\overline{Y}_{GR,2016-2018}$ - the mean of VAT gap as a percentage of VTTL for the post-policy period (2015-2018) in Germany

Results

In this section, we introduce our results without interpretation. These figures are explained with the highlight of the key findings, which are then going to be used in the discussion as well as the conclusion for analysis, interpretation, and evaluation.

DOI: 10.36838/IJHSR820.16

Table 1: France. VTTL, VAT revenue, VAT gap, VAT gap as a percentage of VTTL and its changes (expressed in percentage points), 2010-2018, (EUR Million). While the VAT gap worsened by 3.1 pp in the pre-policy period (2010-2014), it declined by 2.4 pp following the introduction of the cash payment limit in 2015, with the overall net change of only 0.1 pp across the full period.

	Before cash payment limit					After cash payment limit			
	2010	2011	2012	2013	2014	2015	2016	2017	2018
VTTL	144,369	147,654	156,992	158,026	165,520	167,521	168,611	173,840	180,406
VAT Revenue	135,875	140,860	142,840	144,618	148,454	151,680	154,490	162,011	167,618
VAT Gap	10,439	9,359	14,771	14,286	17,066	15,841	14,121	11,829	12,788
VAT Gap as a percent of VTTL	7.2	6.3	9.4	9.0	10.3	9.5	8.4	6.8	7.1
VAT Gap change since 2010									0.1
VAT Gap change since 2015									−2.4
VAT Gap change since 2010–2014					3.1				

Table 2: Germany. VTTL, VAT revenue, VAT gap, VAT gap as a percentage of VTTL and its changes (expressed in percentage points), 2010-2018, (EUR Million). The VAT gap declined by 1.8 pp over the full period, with the post-policy reduction of 2.4 since 2015.

	Before cash payment limit					After cash payment limit			
year	2010	2011	2012	2013	2014	2015	2016	2017	2018
VTTL	202,304	216,260	221,946	224,338	229,881	232,507	239,911	248,382	257,207
VAT Revenue	180,213	189,910	194,034	197,005	203,081	211,616	218,779	226,582	235,130
VAT Gap	21,880	26,510	28,199	27,499	26,800	20,891	21,132	21,800	22,077
VAT Gap as a percent of VTTL	10.8	12.3	12.7	12.3	11.7	9.0	8.8	8.8	8.6
VAT Gap change since 2010									−1.8pp
VAT Gap change since 2015									−0.4pp
VAT Gap change since 2010–2014					0.9pp				

In Tables 1 and 2, we reported VAT gaps (all expressed in millions of euros) in France and Germany over the same time periods (before and after the policy implementation). The tables show that both countries experienced a rise in VAT gap in the pre-policy period. The VAT gap in France had a limited net change over the full time period, so the VAT gap as a percentage of VTTL decreased only by 0.1 pp. Germany, however, experienced a noticeable decline in VAT gap after 2014. Its VAT gap as a percentage of VTTL also decreased more significantly compared to France (by 1.8 pp), suggesting a larger decline compared to France.

Table 3: Mean of the VAT gaps as a percent of VTTL (2010-2018) for France and Germany across pre- and post-policy periods in percentage points (pp) showing the Difference-in-Differences estimate of 2.8 pp.

	Pre-policy period	Post-policy period	Difference
	2010–2014	2015–2018	
France	8.4	7.9	−0.5
Germany	12.0	8.8	−3.2
Difference-in-differences			2.8

Table 3 shows the means of VAT gaps as a percent of VTTL over two different time periods: pre-policy (2010-2014) and post-policy (2015-2018). The table shows that the mean VAT gap as a percent of the VTTL in France was 8.4 pp in 2010-2014, but it then declined by 0.5 pp (to 7.9 pp) for the post-policy period (2015-2018). Similarly, in Germany, the mean VAT gap as a percent of the VTTL has declined, but by 3.2 pp (from 12.0 pp in 2010-2014 to 8.8 pp in 2015-2018). The computed difference-in-difference estimate is therefore 2.8 pp, which indicates that the reduction of the VAT gap as a percent of VTTL in France was smaller by 2.8 pp relative to Germany.

To provide statistical inference, we estimate the difference-in-differences model using an OLS regression framework. The coefficient on the interaction term (France × Post) is 2.715 and statistically significant at the 5% level (p = 0.022). This indicates that the reduction in the VAT gap in France was significantly smaller than in Germany.

Table 4: Regression-based Difference-in-Differences estimates with standard errors, showing the interaction term as the Difference-in-Differences estimate of 2.715, significant at the 5% level.

Variable	Coefficient
Constant	11.96***
	(0.498)
France	-3.540***
	(0.705)
Post	-3.185***
	(0.747)
France × Post	2.715**
	(1.057)
Observations = 18	
$R^2 = 0.7323$	

Notes: Standard errors are reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Mean of the VAT gaps as a percent of VTTL (2010-2018) in percentage points (pp)- Robustness check of the Difference-in-Differences estimate excluding 2015 as the reform year, which gives a Difference-in-Differences value of 2.3 pp.

	Pre-policy period	Post-policy period	Difference
	2010–2014	2016–2018	
France	8.4	7.4	−1.0
Germany	12.0	8.7	−3.3
Difference-in-differences			2.3

Table 5, similarly to Table 3, shows the means of VAT gaps as a percentage of VTTL over two different time periods: pre-policy (2010-2014) and post-policy (2016-2018), but excluding 2015 as the year of the reform implementation. This allows us to mitigate the effect of the policy implementation year and revise the earlier obtained results. The values for the pre-policy period of both countries, therefore, remain unchanged. However, the mean value for the post-policy period in France is slightly smaller: 7.4 pp instead of 7.9 pp, which implies a greater decrease (by 1.0 pp instead of 0.5 pp). In Germany, the mean value of the post-policy period did not change significantly either: 8.7 pp instead of 8.8 pp, which leads to a similar difference of 3.3 pp instead of 3.2 pp. The computed difference-in-difference estimate after the robustness check is 2.3 pp, which is close to the earlier computed value of 2.8 pp. This result implies that our results were consistent and not sensitive to the 2015 shock due to the policy implementation.

Table 6: Size of Shadow economy in France and Germany as a percent of official GDP, 2010 - 2018, from "Taxation of the Informal Economy in the EU."[16] The shadow economy of France has increased over the time period, while in Germany it has decreased.

	Size of Shadow Economy as a percent of official GDP								
	2010	2011	2012	2013	2014	2015	2016	2017	2018
France	11.3	11.0	10.8	9.9	10.8	12.3	12.6	12.8	12.5
Germany	13.5	12.7	12.5	12.1	11.6	11.2	10.8	10.4	9.7

We also examine the relationship between the VAT gap and the shadow economy. Table 6 represents the size of shadow economies in each country as a percentage of its GDP over the same time period. The results for France show an inverse relationship, suggesting that the VAT gap may not fully capture all aspects of informality. For Germany, however, the results show a proportional trend.

The averages, however, do not determine the overall trend accurately, which is why we plotted our data on the following graphs:

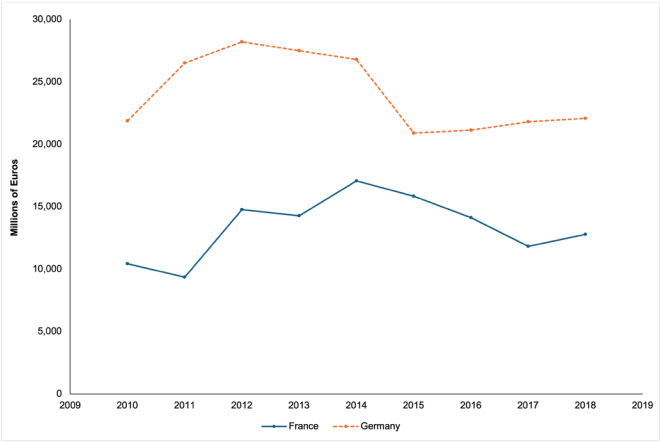


Figure 1: The VAT gap in France and Germany (2010-2018) shows a not perfectly parallel trend. Both countries show a rising VAT gap before 2015, followed by a decline in the post-policy period, with Germany exhibiting a steeper absolute reduction. The author's calculations are based on Poniatowski (2016, 2020).[2,3]

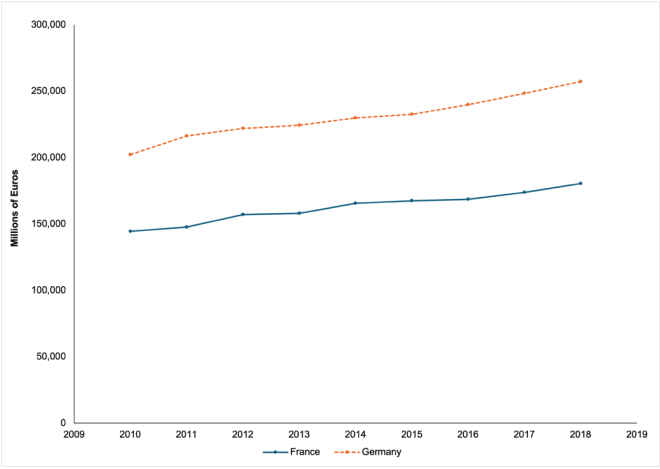


Figure 2: The VAT gap in France and Germany (2010-2018) shows a not perfectly parallel trend. Both countries show a rising VAT gap before 2015, followed by a decline in the post-policy period, with Germany exhibiting a steeper absolute reduction. The author's calculations are based on Poniatowski (2016, 2020).[2,3]

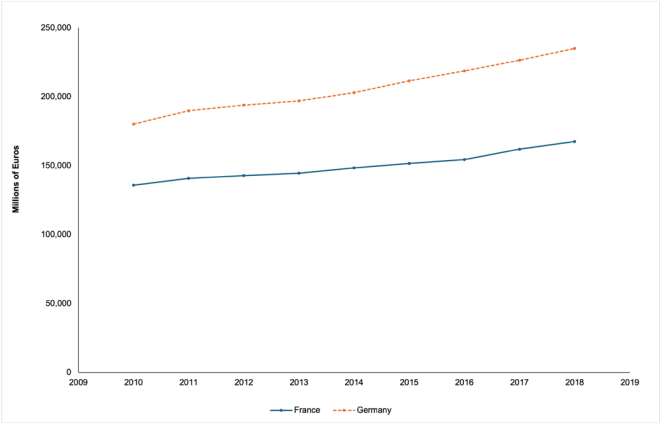


Figure 3: VAT revenue in France and Germany (2010-2018) shows a parallel trend in the two countries. Germany's VAT revenue levels are higher throughout the period, consistent with its larger VTTL base. The author's calculations are based on Poniatowski (2016, 2020).[2,3]

Figures 2 and 3 show increases in VTTL and VAT revenue in both countries, reflecting higher economic activity over time. Germany exhibits higher levels in both variables, consistent with its larger economic size.

Figure 1, nevertheless, shows different patterns in VAT gap behavior of each economy. The dynamics of the curve of France suggest an increasing trend before 2015, which then declines after the cash payment limit implementation. The VAT gap appears to be more stable in France after 2015. In Germany, we found the net VAT gap to be larger due to the greater size of its economy as a whole. Figure 1 shows a significant VAT gap decrease from 2014 to 2015 in Germany, yet this shock is followed by an increasing trend right after. In 2012, the gap between the VAT gap curves of the two countries was at its maximum, and in 2015, it reached its minimum due to the VAT gap decline in Germany, but an overall increasing trend in France. We observed an increase in VAT gap in France and Germany during the pre-policy period and a decrease during the post-policy period.

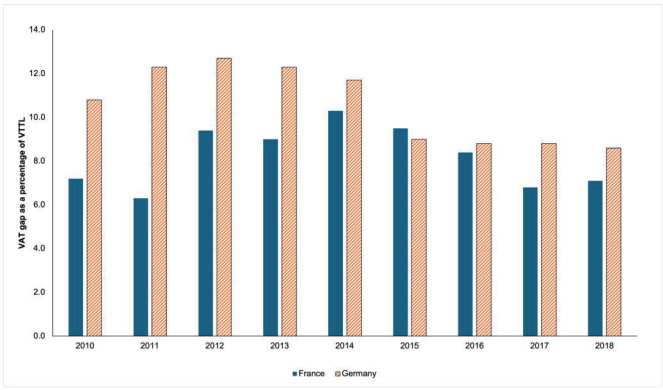


Figure 4: VAT gap as a percentage of VTTL in France and Germany (2010-2018). Germany shows a larger post-policy decline in the VAT gap as a percentage of VTTL, while France displays a more modest reduction following 2015. The author's calculations are based on Poniatowski (2016, 2020).[2,3]

Figure 4 illustrates the VAT gap as a percentage of VTTL in France and Germany. These results were essential for understanding the difference and changes in the VAT gap behavior in each economy with respect to their scales.

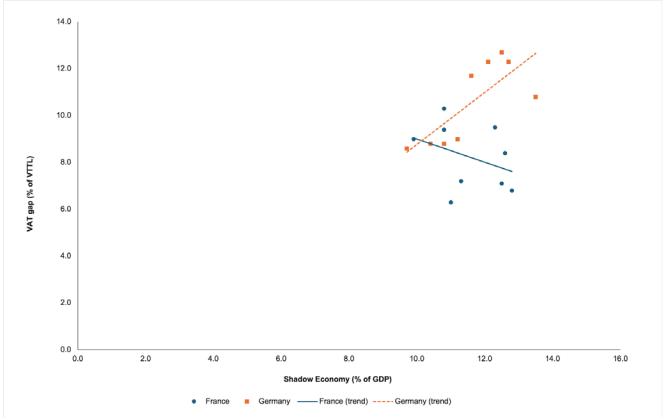


Figure 5: Relationship between VAT gap and Shadow Economy in France and Germany (2010-2018), showing a contrasting correlation pattern between the two countries. France has an inverse relationship between the two variables, while Germany displays a proportional trend. Author's calculations based on Schneider (2022).[16]

DOI: 10.36838/IJHSR820.16

Finally, Figure 5 illustrates the correlation between the shadow economies of France and Germany as percentages of their official GDPs and VAT gap as percentages of their VTTLs. The relationship between the two variables is different in the two countries. In Germany, the shadow economy is positively correlated with the VAT gap, which implies that as the shadow economy increases, the VAT gap increases too. In France, on the other hand, the opposite trend was observed; as the shadow economy increases, the VAT gap decreases. Figure 5 helps us to examine the relationship between the VAT gap and the shadow economy in France and Germany.

Discussion

This section evaluates and discusses whether the 2015 cash payment limit in France has affected its VAT compliance with respect to the results obtained. We find that our hypothesis was not fully supported by the results.

First, while the VAT gap declined in France, the reduction was smaller than in Germany, indicating a weaker improvement. Figure 1 shows that both countries experienced a decline in the VAT gap after 2015. In France, the decline is gradual and includes a small increase in 2018. In contrast, Germany shows a more pronounced drop around 2015, followed by relatively stable levels. This is consistent with the average changes reported in Tables 3, 4, and 5, where the VAT gap decreases by about 0.5 percentage points in France and by about 3.2 percentage points in Germany. Overall, the reduction in the VAT gap is greater in Germany than in France.

Second, Figure 4 allows us to execute the difference-in-difference analysis of the cash payment limit efficiency in France, comparing its VAT gap as a percent of VTTL to the German one, excluding the influence of the two economies' scales (or in other words, sizes). In Figure 4, before 2015, we observed an increase in VAT gap (as a percentage of VTTL), yet a decline right after, which is again consistent with an improvement. In Germany, after 2015, the trend appears to be rather stable.

Third, Tables 3 and 4 show the difference-in-difference estimates that serve as counter-arguments for our hypothesis. The results imply a decline in the VAT gap as a percent of VTTL in both countries; however, with Germany experiencing a major improvement relative to France. The difference-in-difference estimate shows that the VAT gap decrease in France was 2.8 pp (or 2.3 pp) less relative to Germany. Our regression model showed that this evidence is statistically significant at the 5% level. Our results are consistent with the estimate being close to the computed value after the robustness check (2.8 and 2.3 pp, respectively). While the results obtained present a decline of the VAT gap in France after the policy implementation, this decline is not as strong as in Germany, where no such policy was introduced in the same year. So, it is likely that there are other factors that might have influenced the results. For instance, the greater improvement in VAT compliance in Germany might be caused by a more effective tax administration, which eliminates opportunities for tax evasion more than in France. Another possible reason is the difference in the economic structure of the two countries. Even though our method mitigated the influence of the difference in sizes of each economy, it is still impossible to fully eliminate the influence of other differences in institutional, business, and administrative frameworks. Finally, the weaker improvement in the VAT gap reduction of France might suggest that the policy needs a longer period to fully materialize. Additionally, our estimate relies on the assumption that both European economies would've had a similar behavior if no policy took place, which means that it highly depends on how parallel the VAT gap as a percentage of VTTL trends in both countries before 2015 policy implementation are. This is the parallel trends assumption, which was earlier mentioned in the methodology section. Because France and Germany show different trends of the VAT gap as a percentage of VTTL, they are not perfectly parallel, so the results should be interpreted cautiously.

Finally, linking the VAT gap trends to the shadow economy, we found an interesting relationship in Figure 5. In Germany, we observed an evident result of the shadow economy being proportional to the VAT gap. The relationship in France, however, shows the opposite; as the shadow economy increases, its VAT gap decreases. This inverse relationship may have appeared due to multiple factors. Primarily, the dataset used for this paper was relatively small, which suggests its sensitivity to changes. Secondly, the VAT gap can capture only a particular part of the informal sector of an individual economy, as it reflects a tax evasion problem. It does not, however, correlate with all illegal economic activities. That is why even when decreasing the VAT gap, the overall shadow economy levels may increase due to external and other influencing factors that appeared to be greater in magnitude than the VAT gap improvement. Finally, we assume that the specifics of the time period may have also played a role. Therefore, the VAT gap cannot be taken as a perfect proxy for reflecting the shadow economy. From this, we consider the VAT gap to be one of the indicators of the shadow economy, but not its major determinant. Thus, in the short run, this inconsistency does not necessarily imply a negative trend but shows how informality is affected by various aspects.

Conclusion

In this final section of the paper, we conclude our findings as well as indicate the limitations and improvements that could be made.

This paper studied the effect of the 2015 cash payment limit reduction to €1000 in France on its VAT gap. For our analysis, we decided to use the difference-in-difference method, exploring the behavior of such variables as VAT gap, VTTL, VAT revenues, and informality in two strong European economies, France and Germany, over the same period of time (2010-2018).

Our key findings show that our hypothesis was not fully supported by the evidence. Primarily, the VAT gap trends observed in both France and Germany imply a tendency toward improvement. When excluding the scale of each economy, we still observed a decrease in the VAT gap in both countries, with a larger decline in Germany, followed by relatively stable levels

in the post-policy period. However, the difference-in-difference estimate showed a smaller reduction of the VAT gap in France relative to Germany. This result appeared to be statistically significant at a 5% level after running a regression model, and it is consistent, as shown by the robustness check. Finally, we analyzed how the VAT gap changes links to the shadow economies of these two countries, determining an unexpected inverse relationship between these two variables in France. The difference-in-difference result does not support a strong positive policy effect. The findings suggest that the reform had a limited influence on the VAT gap reduction in France, and it is not enough to combat the growth of the shadow economy. In order to fight the informal sector more effectively, this policy should be followed by other government interventions or replaced with other measures. It therefore implies that further research is needed.

Our paper is important not only for economic literature but also for future policy decision-making. By expanding the VAT gap studies with a specific case study of the French economy, we provide an evaluation of the association between the cash payment limit policy and VAT gap improvements, which are currently highly debated. Our results suggest that the policy had a limited effect on reducing the VAT gap. Therefore, the cash payment restrictions alone may not be enough to reduce the VAT gap. Policymakers may need to combine this policy with other measures to improve tax enforcement.

Our method, however, has limitations, as the economies compared are similar to some extent, yet are not identical. The accuracy of the difference-in-difference method could be argued due to the absence of full control of the German economy (which is taken as our control group). Furthermore, there are external factors affecting both the French and German shadow economy, which could influence our results. Consequently, further research is needed to determine the potential impact of cash payment limits outside of France. These factors leave our paper open to future improvement in accuracy and studies.

Acknowledgments

The author wishes to thank her mentor, Ceyhun Elgin, who was supportive providing his helpful guidance throughout the research process.

References

1. OECD. Shining Light on the Shadow Economy: Opportunities and Threats. OECD Publishing, Paris, **2017**. https://www.oecd.org/content/dam/oecd/en/publications/reports/2017/09/shining-light-on-the-shadow-economy-opportunities-and-threats_a9a92285/e0a5771f-en.pdf
2. Poniatowski, G.; Bonch-Osmolovskiy, M.; Belkindas, M. Study and Reports on the VAT Gap in the EU-28 Member States: 2016 Final Report. CASE Research Paper No. 483, CASE – Center for Social and Economic Research, **2016**. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2847658
3. Poniatowski, G.; Śmietanka, A.; Bonch-Osmolovskiy, M. Study and Reports on the VAT Gap in the EU-28 Member States: 2020 Final Report. CASE Research Paper No. 503, CASE – Center for Social and Economic Research, **2020**. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3744157
4. Kalyva, A.; Naudts, H.; Princen, S. The French VAT System and Revenue Efficiency. European Economy Economic Brief 015, European Commission, Directorate-General for Economic and Financial Affairs, **2016**. https://economy-finance.ec.europa.eu/document/download/cdb1539b-46b4-4861-8023-e84d286925dd_en?filename=eb015_en.pdf&prefLang=hu
5. Barbone, L.; Belkindas, M.; Bettendorf, L.; Bird, R. M.; Bonch-Osmolovskiy, M.; Smart, M. Study to Quantify and Analyze the VAT Gap in the EU-27 Member States. CASE Research Paper No. 116, CASE – Center for Social and Economic Research, **2013**. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2365189
6. Passas, N. Report on the debate regarding EU cash payment limitations. Journal of Financial Crime 25 (1), 5–27, **2018**. https://www.emerald.com/jfc/article/25/1/5/222523
7. Zídková, H. Determinants of VAT Gap in the EU. Prague Economic Papers 4, 515–531, 2014. https://pep.vse.cz/pdfs/pep/2014/04/06.pdf
8. Devigne, L.; Laplace, E.; Olivieri-Battestini, J.-T.; de Pastor, R. Access to Cash Remains First-Rate in France. Bulletin de la Banque de France 243/5, Banque de France, **2022**. https://publications.banque-france.fr/sites/default/files/medias/documents/bdf243-5_en_especes.pdf
9. Bouhdaoui, Y.; Bounie, D. Modeling the Share of Cash Payments in the Economy: An Application to France. SSRN Working Paper, **2012**. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2132449
10. Karatairis, C. Understanding the VAT Gap in the EU: Toward Smarter Tax Policies. Business & Entrepreneurship Journal 14 (1), 13–25, 2025. https://www.scienpress.com/Upload/BEJ/Vol%2014_1_2.pdf 11. Rainone, E. Tax Evasion Policies and the Demand for Cash. SSRN Working Paper, **2021**. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3988679
12. Di Caro, P.; Sacchi, A. The Heterogeneous Effects of Labor Informality on VAT Revenues: Evidence on a Developed Country. Journal of Macroeconomics 63, 103190, **2020**. https://www.sciencedirect.com/science/article/abs/pii/S0164070419301818
13. Cuceu, I. C.; Florescu, D.-R.; Văidean, V. L. The Determinants of the VAT Gap. Journal of Risk and Financial Management 25 (2), 277-292, **2024**. https://www.emerald.com/jrf/article-abstract/25/2/277/1226135/The-determinants-of-compliance-VAT-gap?redirectedFrom=fulltext
14. Cléaud, G.; de Castro Fernández, F.; Durán Laguna, J.; Granelli, L.; Hallet, M.; Jaubertie, A.; Maravall Rodriguez, C.; Ognyanova, D.; Palvolgyi, B.; Tsalinski, T.; Weißschädel, K.-Y.; Ziemendorff, J. Cruising at Different Speeds: Similarities and Divergences between the German and the French Economies. European Economy Discussion Paper 103, European Commission, Directorate-General for Economic and Financial Affairs, **2019**. https://economy-finance.ec.europa.eu/document/download/6494542a-fc54-4f45-a5ec-fae7a4f02bc5_en?filename=dp103_en.pdf
15. European Commission, Directorate-General for Taxation and Customs Union. Mind the Gap: Challenges and Opportunities for Tax Compliance and Tax Expenditures in the EU; Publications Office of the European Union: Luxembourg, **2025**. https://op.europa.eu/en/publication-detail/-/publication/7bd6a021-d764-11f0-8da2-01aa75ed71a1
16. Schneider, F.; Asllani, A. Taxation of the Informal Economy in the EU. Study for the Economic and Monetary Affairs Subcommittee on Tax Matters (FISC), Policy Department for Economic, Scientific and Quality of Life Policies, European Parliament, PE 734.007, **2022**. https://www.cyprusbarassociation.org/images/Cases/2025/PARARTIMA-XI.pdf

DOI: 10.36838/IJHSR820.16

Author

Maria Sultanova is a junior at The International School of Nice with a strong passion for Economics and Mathematics. Through her research, she aims to explore the effects of financial policies and pursue a degree in Economics in college.

Appendix

Table A1: Raw data - 26-EU member states: VAT receipts, VTTL, VAT Gap as a share of VTTL, VAT Gap as a share of GDP.[5]

	VAT receipts	VTTL	VAT Gap	VAT Gap as share of VTTL (%)		VAT Gap as share of GDP (%)	
Member State	2011	2011	2011	2011	2000–2011	2011	2000–2011
Austria	23,447	26,915	3,468	13	11	1.2	1.0
Belgium	26,021	30,991	4,970	16	13	1.3	1.1
Bulgaria	3,352	3,956	604	15	16	1.6	1.8
Czech Republic	10,994	15,235	4,241	28	23	2.7	2.1
Denmark	23,869	26,436	2,566	10	10	1.1	1.1
Estonia	1,363	1,664	301	18	15	1.9	1.6
Finland	16,915	19,746	2,831	14	13	1.5	1.3
France	140,506	172,739	32,233	19	16	1.6	1.3
Germany	189,920	216,830	26,910	12	13	1.0	1.0
Greece	15,027	24,790	9,763	39	30	4.7	3.0
Hungary	8,516	12,216	3,700	30	26	3.7	3.0
Ireland	9,782	10,890	1,108	10	8	0.7	0.6
Italy	98,557	134,691	36,134	27	26	2.3	2.1
Latvia	1,368	2,322	954	41	24	4.7	2.3
Lithuania	2,444	3,795	1,352	36	35	4.4	3.9
Luxembourg	2,690	3,242	551	17	12	1.3	0.8
Malta	520	541	21	4	13	0.3	1.0
Netherlands	41,610	45,622	4,012	9	5	0.7	0.4
Poland	29,843	35,253	5,410	15	13	1.5	1.1
Portugal	14,235	16,999	2,764	16	9	1.6	0.8
Romania	11,412	21,760	10,348	48	42	7.9	5.4
Slovakia	4,711	7,484	2,773	37	29	4.0	2.9
Slovenia	3,049	3,375	326	10	7	0.9	0.6
Spain	56,547	71,744	15,197	21	12	1.4	0.8
Sweden	36,610	37,542	932	2	4	0.2	0.4
United Kingdom	130,577	150,064	19,487	13	12	1.1	1.0
EU-26, total	903,884	1,096,841	192,957	18	15	1.5	1.2
EU-26, average				20	17	2.1	1.6

Table A2: Raw data: France. VTTL, VAT revenue, VAT gap, VAT gap as a percent of VTTL and its changes (expressed in percentage points), 2010-2018, (EUR Million).[2,3]

Year	2010	2011	2012	2013	2014	2015	2016	2017	2018	
VTTL	150,550	153,975	163,713	164,791	172,606	165,520	167,521	168,611	173,840	180,406
VAT Revenue	135,578	140,552	142,527	144,301	148,129	148,454	151,680	154,490	162,011	167,618
VAT Gap	14,972	13,423	21,186	20,490	24,477	17,066	15,841	14,121	11,829	12,788
VAT Gap as a percent of VTTL	10.0	9.0	13.0	12.0	14.0	10.3	9.5	8.4	6.8	7.1
VAT Gap change since 2010					4 pp					-3.2 pp

Table A3: Raw data: Germany. VTTL, VAT revenue, VAT gap, VAT gap as a percent of VTTL and its changes (expressed in percentage points), 2010-2018, (EUR Million).[2,3]

Year	2010	2011	2012	2013	2014	2015	2016	2017	2018	
VTTL	199,390	213,145	218,749	221,107	226,570	229,881	232,507	239,911	248,382	257,207
VAT Revenue	180,213	189,910	194,034	197,005	203,081	203,081	211,616	218,779	226,582	235,130
VAT Gap	19,177	23,235	24,715	24,102	23,489	26,800	20,891	21,132	21,800	22,077
VAT Gap as a percent of VTTL	10.0	11.0	11.0	11.0	10.0	11.7	9.0	8.8	8.8	8.6
VAT Gap change since 2010					1 pp					-3.1 pp

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

The Application of Biomimetics to Maritime Transportation Through Antifouling and Hydrodynamics

Max DeWeese

Carlmont High School, 1400 Alameda de la Pulgas, Belmont, CA, 94002, USA; maximiliandeweese@gmail.com

ABSTRACT: Over millions of years of evolution, marine animals have optimized their hydrodynamics and efficiency of movement in water. Elements of marine animals can be replicated to solve real-world design problems. Hydrodynamic features of these animals can be copied, in particular, to help improve the efficiency and emissions of maritime transportation. This paper introduces the growing field of biomimetics and how this engineering tool can be utilized to increase the efficiency of maritime transportation. It is hypothesized that by utilizing biomimetics to replicate aspects of animal biology, the efficiency of maritime transportation can be improved while lowering emissions. The current and future uses of biomimetics in maritime transportation for drag reduction and antifouling are to be compared to current non-biomimetic solutions through a cost-benefit analysis.

KEYWORDS: Engineering, Maritime Transportation, Biomimetics, Sharkskin Denticles, Antifouling.

Introduction

Maritime transportation has played a central role in human society for millennia and remains essential to global trade today. While early civilizations such as the ancient Mesopotamians applied hydrodynamic principles for irrigation and water management, modern maritime transportation faces significant challenges.[1] Fuel inefficiency and associated emissions contribute to climate change, disrupting marine ecosystems.[2] Because maritime transportation cannot be eliminated, improving its efficiency and reducing its environmental impact is critical.

Biomimetics, the application of biological processes and structures to engineering design, offers a promising framework for addressing these challenges.[3] Through millions of years of evolution, marine organisms have developed effective solutions to problems directly relevant to maritime systems, including drag reduction and antifouling.[4,5] Leveraging these biologically optimized strategies provides a pathway toward improving the efficiency and sustainability of maritime transportation.

More broadly, biomimetics is part of a well-established family of nature-inspired design approaches that have been successfully applied across engineering, architecture, and materials science. The existence of related frameworks underscores that drawing inspiration from biological systems is not a speculative idea, but a proven and transferable design paradigm, one that can be meaningfully applied to maritime transportation. Figure 1 provides a visual overview of the broader landscape of nature-inspired design approaches, situating biomimetics within a continuum of strategies that draw on biological systems in different ways. For example, related approaches to biomimetics include biofabrication, which incorporates biological materials into manufacturing; biophilia, which integrates nature into design to support human well-being; and bionics, which applies biological principles to engineering and electronic systems.

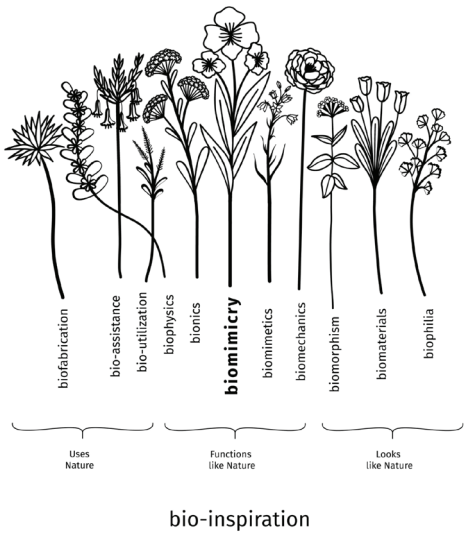


Figure 1: Demonstration of a spectrum of bio-inspired design approaches similar to Biomimetics.[6]

This paper aims to demonstrate the potential of biomimetics in improving the efficiency of maritime transportation. Section 2 focuses on the application of biomimetics in improving the hydrodynamics of maritime transportation vessels to reduce drag, improve efficiency, and reduce emissions. Section 3 provides an overview of historical antifouling strategies and evaluates their performance relative to modern biomimetic antifouling solutions using a cost–benefit framework. Finally, the future of biomimetics itself is discussed briefly in Section 4.

Discussion of Improving the Hydrodynamics of Maritime Transportation Through Biomimetics

Reducing hydrodynamic drag is a key pathway to improving efficiency and reducing environmental impact in maritime transportation.[7] Biomimetics provides valuable insight into

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.17

this problem, as aquatic organisms have evolved effective mechanisms for minimizing drag and enhancing locomotion.

Biological Inspiration: Sharkskin Denticles:

One organism that depends heavily on efficient underwater movement is the shark, for which sharkskin denticles play a key role in enhancing hydrodynamic efficiency. Denticles are small tooth-like structures embedded in sharkskin to reduce drag during swimming. Figure 2 provides a close-up image of these denticles. This has inspired researchers to mimic denticle designs to improve the hydrodynamics of various underwater surfaces.[4]

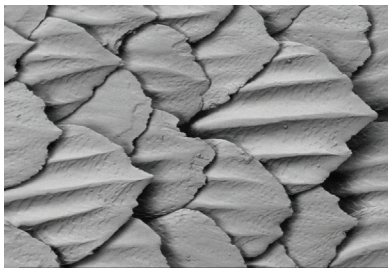


Figure 2: Image of sharkskin denticles.[8]

While sharkskin denticles demonstrate highly effective drag-reduction strategies in biological systems, their performance in engineered applications depends strongly on surface properties that may not be as easily replicated in large maritime structures; yet, the following research studies show great promise for implementing them for the future of maritime transportation.

Role of Surface Flexibility:

Surface flexibility is a critical factor in determining whether biomimetic denticles can effectively reduce hydrodynamic drag. Flexibility influences not only where denticles may be applied on a surface, but also whether their interaction with the surrounding flow can produce performance benefits.

In 2012, Oeffner and Lauder conducted experiments on artificial shark skin, real shark skin, and a flexible Speedo foil material.[4] As part of their testing setup, some denticles were sanded down to form small surface nubs. When applied to rigid surfaces, these denticles showed no consistent improvement in swimming performance. Although a systematic explanation for the lack of performance change was not provided, the surface nub results suggest that denticle surface roughness alone is less important than the interaction between the surface and the flow. In particular, the ability of the shark skin to deform underwater, in addition to the denticles on the surface of the shark skin, provides the performance-enhancing benefits underwater.[4]

In contrast, when surface nub denticles were applied to a flexible surface to create artificial shark skin, swimming speed increased by 12.3%.[4] The testing in this model closely replicated the hydrodynamics of real shark skin underwater. These denticles experienced an average strike number, St between .28 and .38, which is within the range of Strouhal's number of swimming sharks: 0.21 to 0.41.[9] This is defined as $St = fL/U$, where f is the frequency of the foil, L is the characteristic length, and U is the flow speed. Strouhal's number is a dimensionless parameter that expresses the ratio of vibrational-to-flow velocities.[10]

One potential cause behind the improved performance of the denticles on flexible surfaces is due to how a flexible surface might mimic the flesh of a shark. The denticle's improved performance on flexible materials leads towards the potential future use of sharkskin-inspired denticles on athletic swimwear and flexible underwater robots.

To further assess when a surface is flexible enough to benefit from denticles, a scaling analysis can be used. The degree to which a material deforms under fluid loading is characterized by its bending stiffness, B given by:[11]

$$B = \frac{Et^3}{12}(1-\nu^2) \quad (1)$$

where E is Young's modulus, t is the thickness, and ν is Poisson's ratio.[12, 13] The force exerted by the fluid can be estimated using the dynamic pressure, q:[14]

$$q = 1/2\,\rho U^2 \quad (2)$$

Suppose the pressure acts across a length L, the surface area exposed to the flow scales as L^2, where the total bending moment acting on the material is $M=qL^3$. The ratio of the fluid force to the resistance force is then given by $\underline{B}$:

$$\underline{B} = \frac{6(1-\nu^2)\rho U^2 L^3}{Et^3} \quad (3)$$

Equation (3) can be used to predict if a material would benefit from the addition of denticles under certain flow conditions. When $\underline{B} < 1$, the elastic resistance dominates, and the structure does not bend under flow, making it rigid and unlikely to benefit from the addition of denticles. When $\underline{B} > 1$, the fluid forces dominate, and the structure deforms greatly under fluid forces. Under this regime, the structure would benefit most from the addition of denticles.

Several trends can be inferred from equation (3), plots of which are shown in Figure 3 for the case when $\underline{B} > 1$. First, thinner structures benefit more from the addition of denticles. Second, the flow regime is especially important. Below in Figure 3, equation (3) is graphed with $\underline{B}$ as our dependent variable and U as our independent variable. The velocity U ranges from 0 to 15m/s; this ranges from 0 to 29 knots. Figure 3 shows 3 cases at thicknesses of 5 and 25mm. Additionally, a length of 1 meter, a Poisson's ratio of .3 and a Young's modulus of $(5)10^{10}$ is utilized. Water density at approximately 3.98 C is assumed.

The quadratic relationship between $\underline{B}$ and U is evident in all 3 cases; additionally, the thickness of the material significantly impacts the value of $\underline{B}$. At the lowest thicknesses of 5mm, $\underline{B} > 1$ at the majority of speeds; however, a thicker material requires greater speeds for $\underline{B} > 1$. For example, with a thickness of 25 mm, U must exceed 12 m/s for $\underline{B} > 1$. Consequently, at thinner values the use of denticles is almost always beneficial; however, as the thickness increases, the velocity of the ship must exceed greater and greater values for denticles to become beneficial.

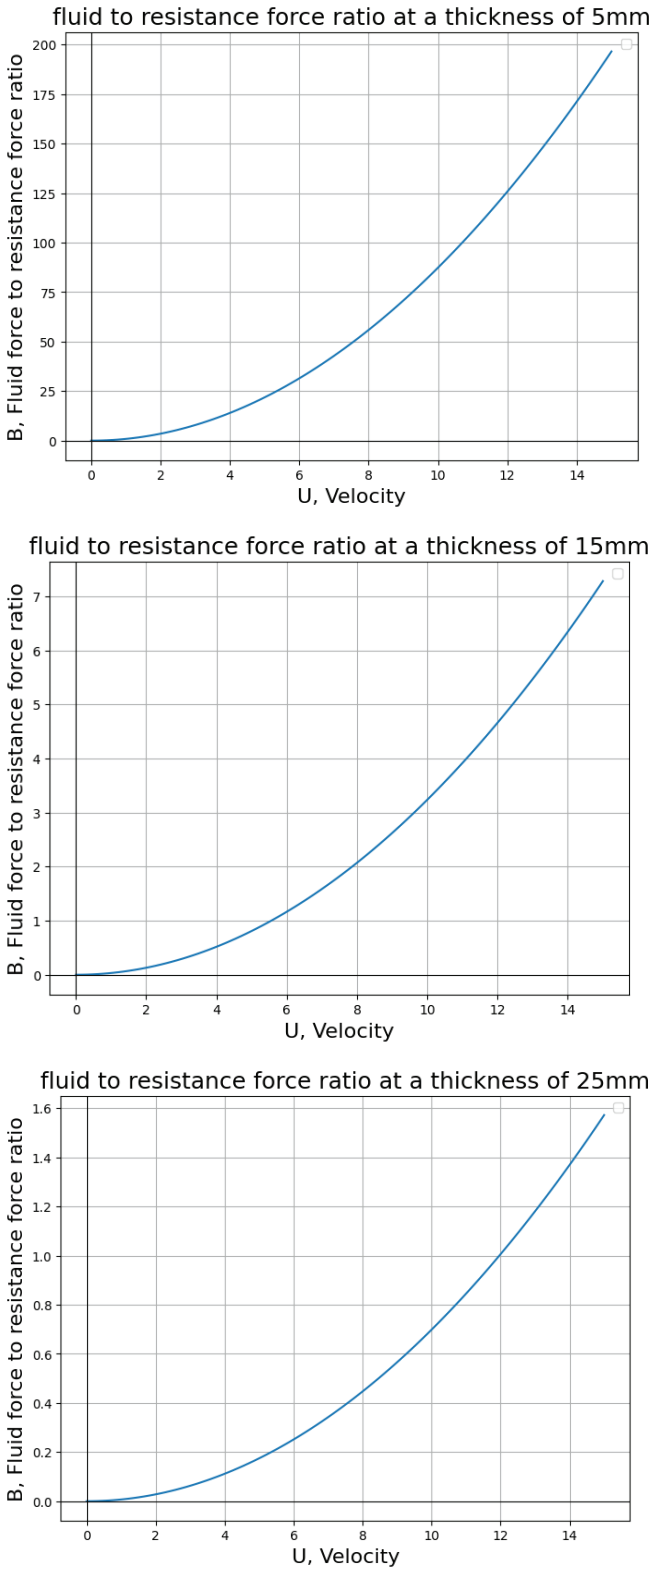


Figure 3: Graphs of equation 3 plotted with thicknesses of 5, 15, and 25 mm. A quadratic relationship exists between the velocity and force ratios. Additionally, the thickness of the material impacts the curve; thicker materials require greater velocities for the denticles to be beneficial. Plots generated by the author.

Manufacturing Constraints and Feasibility:

Although shark skin denticles have the potential to improve the drag of maritime transportation and robotics, they suffer from being hard to manufacture. Due to the tiny size of these denticles and the harsh conditions they endure, the manufacturing and design of these denticles becomes a significant roadblock. In 2024, Graybill and Xu analyzed different manufacturing techniques, including bioreplication, molding, 3D printing, and metal CNC machining for the creation of biomimetic denticles.[7] These denticles mimicked those of the Shortfin Mako Shark.[7]

Molding:

Molded silicone denticles were the closest in size to real shark skin denticles, which were also the smallest of all manufacturing techniques. The molded denticles experienced a drag reduction of 9.7% at 0.9 m/s to 18.6% at 0.45 m/s.[7] Compared to a flat PU sheet, molded denticles experienced 31% more drag reduction, and replicated denticles a 36.7% reduction with MSN texturing, which is close to the nearly 39% of sharkskin.[7]

CNC Machining:

CNC machining these denticles from Aluminum is a complex and traditional process. Denticles were machined with spacings from 0.150 to 1.001 mm and tested with flow velocities ranging from .55 to 4.52 m/s.[7] The smallest design at 0.150mm experienced the greatest average drag reduction at 13.05%, peaking with 27.7% at 2.36 m/s.[7] Along with producing the worst results, CNC-machined denticles struggle with feasibility, as traditional CNC manufacturing could be more expensive than molding or 3D printing.

3D Printing:

Another primary manufacturing method analyzed is 3D printing with thermoplastics. These were tested in flow speeds between .1 to .5 m/s with widths ranging from .15 mm to 3 mm in width.[7] Although 3D printing technology is more time-efficient in outputting the desired product, there are limitations in the context of denticle production, where achieving smaller-sized denticles under 2 mm is much harder. Bigger denticles (~2 mm) experienced reduced drag reduction benefits as the form drag of a larger denticle overcomes its hydrodynamic benefits.[7] A drag reduction of 31% was experienced with the smallest denticle of width 1.53 mm.[7]

One of the potential benefits of 3D printing compared to molding is that they are easier to produce on smaller scales; however, on large scales, molding is a much more scalable option.[7] In 2018, Domel *et al.* further tested 3D printed denticles and found that the smallest denticle with a chordwise length of 2.1 mm reduced the most drag.[15] The denticles were 3D printed out of a thermoplastic with a rubber-like membrane designed to simulate the actual materials of sharks.[15] Static water-tank testing was conducted to measure drag forces, with denticles tested at three different sizes under the hypothesis that smaller denticles would reduce drag relative to larger ones.[15] Additionally, all denticles experienced increased drag at faster flow speeds and similar Reynolds numbers below 2000, suggesting the overall flow is similar between each.[15]

DOI: 10.36838/IJHSR820.17

Overall, all machining methods found that smaller denticles under 2 mm in width most consistently reduced drag.[16]

Implications for Future Maritime Technologies:

The use of sharkskin denticles to reduce hydrodynamic drag and improve efficiency is incredibly promising.[15] Optimizing for reduced drag, smaller denticles of length 1 to 2 mm on flexible surfaces produce the best results. Currently, biomimetics research indicates that denticle-based designs are most easily transferable for applications involving flexible or deformable surfaces, such as competitive swimwear and underwater robotic systems. In underwater robotics, where compliant outer structures are often used to enhance maneuverability and efficiency, the integration of denticle-inspired surface textures may offer tangible performance improvements.

In contrast, the application of denticles to large maritime vessels faces challenges that require more in-lab investigations. Ships and larger vessels suffer from not being able to have more flexible hulls without sacrificing other important components. Additionally, larger aquatic vessels would require more denticles, which could increase the cost of manufacturing, assembling, and maintaining so many small structures on the ship's hull.[17] Drag reduction through denticle-inspired designs does have the potential to lower fuel consumption and associated emissions, but the economic and environmental costs of producing and maintaining such surfaces must be carefully weighed against these benefits.

Despite these challenges, sharkskin denticles illustrate biomimetics' capacity to inspire novel solutions to modern maritime challenges, with recent studies examining denticle geometry, surface compliance, and hydrodynamic performance to evaluate their potential translation from biological systems to engineered marine surfaces, including ship hulls.[17]

■ Discussion Of Applying Antifouling to Biomimetics

Biofouling as an Efficiency and Environmental Problem:

Biofouling, the accumulation of marine life, has always been a problem in the maritime transportation sector. The accumulation of marine life on the hull of a ship results in an increase in weight and friction, causing an approximate 40% increase in fuel consumption or 72 million tonnes each year.[18] Additionally, the overall cost associated with hull biofouling in the Navy alone is estimated to be $56 million per year.[19] Although biofouling can impact fuel emissions and costs, biofouling also contributes to the spread of invasive organisms. This can occur when a biofouled ship travels long distances. In the state of Hawai'i, nearly 80% of non native aquatic species have been introduced through biofouling.[20] This section explains biofouling, a brief history of antifouling solutions, and their pros and cons. Furthermore, current antifouling methods are compared to biomimetic alternatives.

Biofouling begins with the formation of a biofilm on a submerged surface, initiated by the colonization of microorganisms such as diatoms and protozoa. These early-stage, single-celled organisms create a conditioning layer that facilitates the subsequent settlement of algae spores and larger marine organisms. Figure 4 illustrates this biofouling progression over the course of one year.

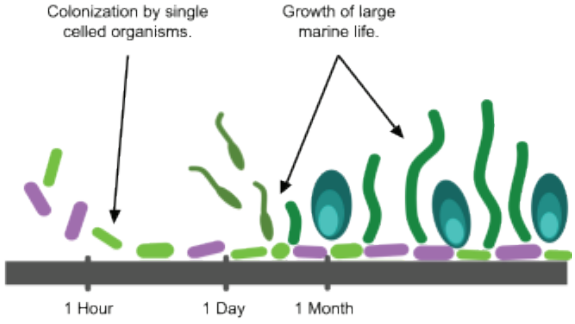


Figure 4: Model of the biofouling process over time. Image created by the author.

With biofouling's impact on overall efficiency and movement, antifouling techniques were guaranteed to arise. Antifouling is a design technique to prevent biofouling or the accumulation of marine life on the hull of ships. One early example of antifouling is the plating of a ship's hull with copper and lead to remove marine organisms.[21] Figure 5 is an example of biofouling on the hull of a ship.



Figure 5: Biofouling on a workboat.[22]

Conventional Antifouling Methods and Their Limitations:

Several modern antifouling solutions include regular cleaning of surfaces and applied coatings or paints. In the 1960's, antifouling paints with Zinc, Mercury, and other toxic compounds were used that had a considerable impact on the marine environment.[23] One of the most common antifouling paint solutions, known as TBT, is a compound that lasts up to 4-5 years.[18] The toxic effects of TBT are well researched and have extreme impacts on marine biology. For example, in 1982, France banned the use of TBT on vessels less than 25 meters in length.[18] Additionally, TBT has been found in human blood and liver.[24] Compared to traditional non-biomimetic antifouling solutions, biomimetic solutions are a lot more promising when it comes to environmental impact.

Biomimetic Antifouling Approaches:s

One emerging biomimetic antifouling solution includes the use of algae-inspired biomimetic haloperoxidases.[5] This solution is based on haloperoxidases that algae naturally secrete to prevent the attachment of microorganisms. Different haloperoxidases are still being tested for use, as a perfect solution has not yet been found, but this remains a promising antifouling solution. Additionally, naturally produced enzymes struggle with long-term stability and cost compared to their artificially produced alternatives. Another emerging biomimetic solution

to antifouling is the use of capsaicin films for their antifouling and antibacterial uses.[25] Multiple films composed of capsaicin-filled nanocapsules were tested for their antifouling capabilities and ability to release capsaicin in a gradual and controlled manner. The films maintained over 90% antibacterial properties after 60 days of performance.[25] Furthermore, as time goes on, biomimetic antifouling solutions will become cheaper and more effective. Environmentally friendly antifouling solutions have gained significant interest over the past couple of years as traditional coatings are banned.[26] As scientific development and research around antifouling progresses, biomimetics will emerge as a promising solution.

Comparative Cost–Benefit Analysis:

Section 3.4. presents a comparison between a conventional antifouling agent, TBT, and a biomimetic alternative based on haloperoxidase mechanisms. Table 1 provides a cost-benefit analysis comparing biomimetic and traditional antifouling strategies, highlighting the advantages of biomimetic solutions. The analysis indicates that while TBT is a cheap, reliable, and effective solution, biomimetic haloperoxidases outshine it. This can be seen in their minimal environmental impacts and potential to be just as reliable and cheap as TBT after further research and development.

Table 1: Cost-benefit analysis comparing TBT and biomimetic haloperoxidases antifouling solutions. While TBT remains a highly effective and economically attractive antifouling agent, its severe environmental and regulatory drawbacks limit its long-term viability. On the other hand, biomimetic haloperoxidases offer substantially lower environmental impacts and the potential for future improvements in cost and commercial feasibility. Table created by the author.

	Benefits	Costs
TBT	TBT is incredibly effective and durable, lasting 4-5 years.[18]	TBT significantly negatively impacts the environment. For example, in 1962, TBT significantly impacted oyster farms in France. 80-100% of Oysters affected by TBT experienced developmental anomalies and shell malformations.[18]
	TBT was invented in the 1950's and is incredibly well tested compared to more innovative antifouling solutions. Although TBT is well tested compared to Haloperoxidases, this is a weakness of all emerging technologies.[27]	TBT is heavily restricted legally, following TBT's environmental impacts. In 2008, the International Maritime Organization banned use of TBT on all vessels.[21]
	TBT is significantly cheaper than newer Biomimetic Antifouling solutions.[27]	TBT's environmental impact reaches far beyond marine biology. Studies have found TBT in humans. Additionally, there is a wide range of potential adverse effects of TBT exposure.[24]
Biomimetic Haloperoxidases	Biomimetic solutions have significantly less environmental impact compared to TBT.[5]	This solution is currently not fully developed compared to a much more accessible and cheaper TBT.[5] As biomimetic haloperoxidases continue to develop, this problem will most likely be minimized.
	Biomimetic haloperoxidases for antifouling solutions have the potential for much higher stability and lower costs.[5]	The current use of naturally produced enzymes for antifouling struggles with high production costs and long-term stability.[5]

Conclusion

The field of biomimetics has immense potential in revolutionizing both the hydrodynamics and antifouling of maritime transportation. As ocean exploration expands, new biological discoveries will continue to drive biomimetic innovations, enabling ongoing improvements in maritime technology both now and in the future. Despite its potential, biomimetics suffers from the complexity of nature. Most organisms and their functions perform multiple roles on many different layers. Consequently, biomimetic solutions might suffer from side effects that aren't yet fully understood. The complexity of reproduction in nature also doesn't mimic modern manufacturing logistics, making many biomimetic solutions challenging logistically.

Although antifouling can improve the sustainability of maritime transportation, traditional antifouling solutions have the opposite result. TBT, a common traditional antifouling solution, has immense negative impacts on the environment, which greatly outweigh the positives of antifouling itself. Fortunately, biomimetic solutions reduce the environmental impacts of TBT, providing an overall boost to the sustainability of maritime transportation. Future research regarding this subject should involve a more in-depth cost-benefit analysis of TBT, Biomimetic Haloperoxidases, and other biomimetic antifouling solutions, such as the use of capsaicin films. Another area where biomimetics improves the sustainability of maritime transportation is the hydrodynamics of maritime vessels. Improved drag can not only reduce fuel consumption but also transportation time, emissions, and costs. The use of shark denticles to reduce the hydrodynamic drag of aquatic vessels is delved into, along with the logistics and physics behind it. Along with antifouling and drag reduction, biomimetics could be applied to other fields within maritime transportation, such as propulsion and communications.

s

Acknowledgments

I would like to thank Dr. Andrea Giusti for his incredibly helpful mentorship during the writing process; additionally, I want to thank my parents for their unwavering encouragement and support.

References

1. Rost, S. Navigating the Ancient Tigris – Insights into Water Management in an Early State. *Journal of Anthropological Archaeology* **2019**, *54*, 31–47. https://doi.org/10.1016/j.jaa.2019.01.005.
2. Deng, S.; Mi, Z. A Review on Carbon Emissions of Global Shipping. *Mar Dev* **2023**, *1* (1), 4. https://doi.org/10.1007/s44312-023-00001-2.
3. Thanigaiarasu, P. Biomimetics in the Design of Medical Devices. In Trends in Development of Medical Devices; Elsevier, 2020; pp 35–41. https://doi.org/10.1016/B978-0-12-820960-8.00003-4.
4. Oeffner, J.; Lauder, G. V. The Hydrodynamic Function of Shark Skin and Two Biomimetic Applications. *Journal of Experimental Biology* **2012**, *215* (5), 785–795. https://doi.org/10.1242/jeb.063040.
5. Yang, M.; Wang, N.; Zhang, R.; Khan, S.; Hou, B.; Sand, W. Biomimetic Haloperoxidases for Antifouling on the Surface of Marine Materials: A Review. *Mater. Adv.* **2025**, *6* (2), 479–489. https://doi.org/10.1039/D4MA00700J.
6. *Image from "What (really) is biomimicry?"* Learn biomimicry. https://www.learnbiomimicry.com/blog/what-is-biomimicry?srslti-d=AfmBOooVuQJukHjJHt4jDpaJ_JGdvvr0S7KRv0DuSkG-3ZOBX1kA51KKe (accessed 2026-02-16).
7. Graybill, M. T.; Xu, N. W. Experimental Studies of Bioinspired Shark Denticles for Drag Reduction. *Integrative And Comparative Biology* **2024**, *64* (3), 742–752. https://doi.org/10.1093/icb/icae086.

DOI: 10.36838/IJHSR820.17

8. *Image from "Biomimicry Shark Denticles."* Smithsonian. https://ocean.si.edu/ocean-life/sharks-rays/biomimicry-shark-denticles (accessed 2026-02-16).
9. Flammang, B. E.; Lauder, G. V.; Troolin, D. R.; Strand, T. Volumetric Imaging of Shark Tail Hydrodynamics Reveals a Three-Dimensional Dual-Ring Vortex Wake Structure. *Proc. R. Soc. B.* **2011**, *278* (1725), 3670–3678. https://doi.org/10.1098/rspb.2011.0489.
10. Katopodes, N. D. Viscous Fluid Flow. In *Free-Surface Flow*; Elsevier, 2019; pp 324–426. https://doi.org/10.1016/B978-0-12-815489-2.00005-8.
11. Morvaridi, M.; Brun, M. Perfectly Matched Layers for Flexural Waves in Kirchhoff–Love Plates. *International Journal of Solids and Structures* **2018**, *134*, 293–303. https://doi.org/10.1016/j.ijsolstr.2017.11.009.
12. Belyadi, H.; Fathi, E.; Belyadi, F. Rock Mechanical Properties and In Situ Stresses. In *Hydraulic Fracturing in Unconventional Reservoirs*; Elsevier, 2017; pp 207–224. https://doi.org/10.1016/B978-0-12-849871-2.00013-7.
13. *Young's Modulus.* University of Washington Materials Science & Engineering Resources. https://depts.washington.edu/matseed/mse_resources/Webpage/Biomaterials/young%27s_modulus.htm (accessed 2026-02-16).
14. Vakil, A.; Green, S. I. Stagnation Pressure, the Bernoulli Equation, and the Steady-Flow Energy Equation. *International Journal of Mechanical Engineering Education* **2011**, *39* (2), 130–138. https://doi.org/10.7227/IJMEE.39.2.4.
15. Domel, A. G.; Domel, G.; Weaver, J. C.; Saadat, M.; Bertoldi, K.; Lauder, G. V. Hydrodynamic Properties of Biomimetic Shark Skin: Effect of Denticle Size and Swimming Speed. *Bioinspir. Biomim.* **2018**, *13* (5), 056014. https://doi.org/10.1088/1748-3190/aad418.
16. Qin, L.; Hafezi, M.; Yang, H.; Dong, G.; Zhang, Y. Constructing a Dual-Function Surface by Microcasting and Nanospraying for Efficient Drag Reduction and Potential Antifouling Capabilities. *Micromachines* **2019**, *10* (7), 490. https://doi.org/10.3390/mi10070490.
17. Fu, Y. F.; Yuan, C. Q.; Bai, X. Q. Marine Drag Reduction of Shark Skin Inspired Riblet Surfaces. *Biosurface and Biotribology* **2017**, *3* (1), 11–24. https://doi.org/10.1016/j.bsbt.2017.02.001.
18. Amara, I.; Miled, W.; Slama, R. B.; Ladhari, N. Antifouling Processes and Toxicity Effects of Antifouling Paints on Marine Environment. A Review. *Environmental Toxicology and Pharmacology* **2018**, *57*, 115–130. https://doi.org/10.1016/j.etap.2017.12.001.
19. Schultz, M. P.; Bendick, J. A.; Holm, E. R.; Hertel, W. M. Economic Impact of Biofouling on a Naval Surface Ship. *Biofouling* **2011**, *27* (1), 87–98. https://doi.org/10.1080/08927014.2010.542809.
20. *Ballast Water/Biofouling*. HAWAI'I BALLAST WATER AND BIOFOULING BIOSECURITY PROGRAM. https://dlnr.hawaii.gov/ais/ballastwaterbiofouling/ (accessed 2026-02-16).
21. Dafforn, K. A.; Lewis, J. A.; Johnston, E. L. Antifouling Strategies: History and Regulation, Ecological Impacts and Mitigation. *Marine Pollution Bulletin* **2011**, *62* (3), 453–465. https://doi.org/10.1016/j.marpolbul.2011.01.012.
22. Reams, B. *Image from "Fouling Control on Workboats."* MarineLink. https://www.marinelink.com/news/fouling-control-workboats-441830 (accessed 2026-02-16).
23. Romeu, M. J.; Mergulhão, F. Development of Antifouling Strategies for Marine Applications. *Microorganisms* **2023**, *11* (6), 1568. https://doi.org/10.3390/microorganisms11061568.
24. Antizar-Ladislao, B. Environmental Levels, Toxicity and Human Exposure to Tributyltin (TBT)-Contaminated Marine Environment. A Review. *Environment International* **2008**, *34* (2), 292–308. https://doi.org/10.1016/j.envint.2007.09.005.
25. Hao, X.; Chen, S.; Qin, D.; Zhang, M.; Li, W.; Fan, J.; Wang, C.; Dong, M.; Zhang, J.; Cheng, F.; Guo, Z. Antifouling and Antibacterial Behaviors of Capsaicin-Based pH Responsive Smart Coatings in Marine Environments. *Materials Science and Engineering: C* **2020**, *108*, 110361. https://doi.org/10.1016/j.msec.2019.110361.
26. Jin, H.; Tian, L.; Bing, W.; Zhao, J.; Ren, L. Bioinspired Marine Antifouling Coatings: Status, Prospects, and Future. *Progress in Materials Science* **2022**, *124*, 100889. https://doi.org/10.1016/j.pmatsci.2021.100889.
27. Liu, D.; Shu, H.; Zhou, J.; Bai, X.; Cao, P. Research Progress on New Environmentally Friendly Antifouling Coatings in Marine Settings: A Review. *Biomimetics* **2023**, *8* (2), 200. https://doi.org/10.3390/biomimetics8020200.

Author

Max DeWeese is a junior at Carlmont High School, Belmont, California. He is interested in conducting research in mechanical engineering and robotics during his free time.

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Motorsport Brake Systems: Numerical Simulation and Analysis of Heat Distribution in Perforated Brake Discs

Max B. Hong

ACG Parnell College, 9 Davis Crescent, Auckland, 1052, New Zealand; maxbzhong@gmail.com

ABSTRACT: In motorsport, the cooling of the braking system is very important. Sustained high temperatures, resulting from repeated braking under high loads or at high speeds, cause a highly undesirable temporary and sudden reduction in braking power. To reduce these effects, brake discs and pads in racecars are typically made up of Carbon-carbon composite Materials; however, they too can experience brake fade, resulting in decreased braking performance, with sustained high temperatures leading to structural damage to the brake disc, warped brake discs due to uneven thermal expansion, and the boiling of the brake fluid causing the brake to feel spongy and unresponsive. Additionally, at these high temperatures, the brake disc surface may oxidize, causing the dust from the friction of the brake disc as well as surrounding components to possibly ignite at temperatures beyond 1200 °C, causing a fire, thus preventing the racecar from continuing to race. This research investigates the cooling of perforated and non-perforated brake discs by numerically simulating the heat equation in one, two, and three dimensions, using convective boundary conditions at the edges of the simulated object (including at its boundaries). Results indicate that geometry, specifically the surface area of the brake disc due to the number of perforations, strongly influences cooling rates.

KEYWORDS: Mechanical Engineering, Heat Equation Modeling, Finite Difference Method, Convective (Robin) Boundary Conditions, Brake Disc Thermal Simulation.

Introduction

In motorsport, braking systems are subjected to extreme thermal loading, with rapid, repeated high-speed deceleration events generating substantial frictional heat at the brake pad–brake disc interface.[1] This heat generation occurs consistently throughout long race durations, leading to continuous cycles of heating and cooling. Therefore, effective heat dissipation is essential for maintaining braking performance and preventing brake fade, a phenomenon in which braking torque is reduced at elevated temperatures due to temperature-dependent changes in the frictional behavior of the brake disc–brake pad interface (and the brake disc surface, which may become glazed or smooth due to friction), which reduces braking performance.[2] As such, thermal behavior is a critical engineering consideration in motorsport, motivating detailed analysis of how braking systems dissipate heat under sustained loading.

In this research, carbon-carbon composite brake materials were used in the modeling due to their persistent application in most international racing disciplines, such as Formula 1, WEC (including LMH, LMDh and LMGT3), and Indycar,[3,4] and their material properties resulting from its structure (carbon fibers embedded in a carbon matrix, hence the name):[5] high thermal diffusivity, low density, and superior resistance to thermal degradation (compared to traditional metallic-based brake systems and carbon ceramic brake systems found in high-end sports cars). These properties allow carbon-carbon composite components to withstand repeated temperature cycling (repeated cooling to low temperatures and heating to extremely high temperatures) more effectively than metallic-based systems.[6] As such, highlighting why carbon-carbon composite braking material properties were chosen to be used in these modeled simulations.

Mathematical modeling provides a meticulous foundation for understanding how heat diffuses (propagates) through brake pads and discs under real-world racing conditions. The governing model used throughout this study is the heat equation, because it directly represents the fundamental physical mechanism of thermal diffusion: differences in temperature within a solid drive conduction from hotter regions to cooler regions.[7] This is the same behavior expected in brake pads and discs after frictional heating during braking introduces thermal energy at the surface.[8] Once heat enters the brake discs (and brake pads), the internal redistribution of temperature is dictated by conduction. The heat equation provides a direct and established way to represent that redistribution over time and space, also allowing the spatial intensity of heat (hotspots and gradients) to be predicted over time. For these simulations, the heat equation also provides a malleable foundation for testing how geometrical conditions influence temperature profiles, essential when modeling perforated brake discs in motorsport braking systems.

This paper explicitly considers the heat equation in one, two, and three dimensions.[9]

The one-dimensional form serves as a baseline conduction case, useful for validating numerical behavior when heat flow is approximated along a single direction in space:

$$\frac{\partial u}{\partial t} = \alpha \frac{\partial^2 u}{\partial x^2}$$

where $u(x,t)$ is the temperature along the rod, (dependent on position x and time t) and α is the thermal diffusivity of the material (Carbon-carbon composite). The modeling is then

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.18

extended to two dimensions to capture temperature variation across a planar domain:

$$\frac{\partial u}{\partial t} = \alpha\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right)$$

where $u(x,y,t)$ is the temperature along the rod and α is the thermal diffusivity of the material (Carbon-carbon composite).

The three-dimensional form is then included to account for conduction in the depth direction:

$$\frac{\partial u}{\partial t} = \alpha\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2}\right)$$

where $u(x,y,z,t)$ is the temperature along the rod, and α is the thermal diffusivity of the material (Carbon-carbon composite).

NOTE: α is calculated by the equation $\alpha = \frac{k}{\rho C_p}$ where k is the thermal conductivity, ρ is the density, and C_p is the specific heat capacity.

Because brake discs are circular in geometry (flat cylindrical shape), the study further established the transformation from Cartesian coordinates (x,y,z) to cylindrical coordinates (r,θ,z) using the coordinate relations $x = r\cos\theta$, $y = r\sin\theta$, and $z = z$. By applying the multivariable chain rule, the Laplacian operator becomes:

$$\nabla^2 u = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial u}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 u}{\partial \theta^2} + \frac{\partial^2 u}{\partial z^2}$$

Boundary conditions play a critical role in predicting brake-disc temperatures, as the heat equation only describes how temperature diffuses within the brake disc itself, not specifying how the brake disc interacts with its surroundings. In motorsport applications, heat is continually removed at the boundaries: exposed faces (and edges) exchange heat with the surrounding fluid, and hollow structures within the disc (such as perforations in the disc) act as additional cooling boundaries. Thus, in order to produce a realistic simulation of the brake disc temperature distribution, it is imperative that the boundary conditions are assigned.[10] Three boundary conditions were analyzed. Dirichlet (fixed temperature) boundary conditions assume the boundary is held at a constant temperature. This provides a theoretical lower bound on temperature but overestimates cooling compared with real brake discs, whose surfaces heat up significantly under load. Neumann boundary conditions assume zero heat flux across the boundary; thus, heat cannot leave the system. While this may be useful for testing pure conduction behavior or establishing an upper bound for temperature retention, it is not representative of braking systems in airflow because the parts of the braking system (notably the brake discs and brake pads) always lose heat to their surroundings. By contrast, convective (Robin) boundary conditions model heat loss in a way that closely matches real braking behavior: Brake discs cool primarily through convection to air, especially in motorsport, where airflow is manipulated and forced, and ventilation features are designed to accelerate the rate of heat removal. Convective conditions accurately reproduce this by linking heat flux to the temperature difference between the surface and the ambient environment, allowing the model to represent thermally active boundaries, including surface and internal circular openings, that remove heat at a rate dependent on how hot the material is. Because this aligns with how brake components exchange heat with airflow and cooling channels, convective boundary conditions are the most physically realistic of the three. The convective (Robin) boundary condition is written as:[11]

$$-k\frac{\partial u}{\partial n} = h(u - T_\infty)$$

where $\frac{\partial u}{\partial n}$ represents the outward normal temperature gradient at a surface. This normal derivative is physically meaningful both at the outer edges of the disc/plate and at the interior boundaries of the circular perforations for cooling.

To calculate solutions for realistic brake geometries, the heat equation was solved numerically using the finite difference method; spatial and temporal derivatives were approximated using an FTCS (Forward-Time Central-Space) discretization, and the resulting discrete system was advanced in time to obtain the temperature field.

Using this modeling framework, the brake disc is modeled as a 0.5 *m* x 0.5 *m* two-dimensional square plate in four configurations: no perforations, one 10 cm diameter circular perforation, 16 evenly spaced 1 cm perforations, or 25 evenly spaced 1 cm perforations. Convective boundary conditions are imposed at the perforation (hole) boundaries, promoting internal cooling, and at the outer edges, ambient cooling, and temperature fields are sampled at regular intervals over 1000 *s*. This will be expanded on in the methodology section.

The goal of this research is to analyze brake-disc thermal behavior through mathematical modeling of heat distribution under geometric constraints, focusing on how surface-area changes affect heat dissipation pathways. Notably, the perforations are not treated as a cosmetic feature but rather as a quantifiable geometric intervention: increasing the internal boundary length (and therefore effective cooling area) changes the temperature field by introducing additional locations where heat can be transferred to the surrounding fluid. To isolate this effect, a carbon-carbon composite brake disc with dimensions similar to those of a 2014 Mercedes W05 F1 car is modeled under 2 controlled configurations:

1) 1000 perforations of diameter 3 *mm*
2) 800 perforations of diameter 3 *mm*

With convective (Robin) boundary conditions applied to the entire surface exposed to the surrounding fluid. Outputs are generated at regular time intervals.

The 2014 Mercedes W05 F1 Car was chosen as it represents the turning point of Formula 1, bringing the historical motorsport into the modern era with its advanced hybrid power unit, molded by the regulations of that year.

This aim is motivated by the context that motorsport brakes undergo extreme, rapidly changing thermal loads, where large thermal gradients and high temperatures can induce brake fade (loss of braking power due to the excessively high temperature of the brakes) and where accurate heat-transfer understanding becomes a priority. Existing work discussed in this project includes analytical approaches (i.e., moving heat source simulations), experimental approaches (i.e., thermographic imaging), and FEM simulations in platforms such as ANSYS.[1,12] However, the literature considered in this research

often does not closely examine brake-disc thermal interactions using simplified numerical solvers in a way that makes geometry-boundary-temperature relationships easy to interpret

Consequently, this research differs from several prior approaches in two main ways. First, it deliberately prioritizes a finite-difference framework, discretizing the (one, two, and) three-dimensional heat equation using an explicit FTCS formulation such that the link between the governing partial differential equation and the resulting temperature fields remains explicit. In the present work, interpretation is reported primarily on the central cross-section of the brake-disc domain to make spatial trends directly comparable across time. Second, rather than treating convection only on the external boundary, the study explicitly applies convective (Robin) cooling at internal perforation boundaries, as at the external boundaries, such that surface-area changes are evaluated through a physically meaningful mechanism, heat exchange with surrounding fluid, at every cooling interface.

Ultimately, the contribution of this study is a controlled numerical investigation, grounded in the heat equation and FTCS discretization, into how brake-disc relevant geometry reshapes volumetric heat distribution and cooling trends. By comparing different configurations of the brake disc under consistent boundary modeling, this work provides an interpretable basis for connecting geometric design choices to heat dissipation outcomes in carbon-carbon composite brake-disc analogs.

Methods

The finite difference method (FDM) was chosen in preference to the Euler method because the FDM is a complete and consistent numerical discretization of the heat equation,[9] whereas the Euler method is only a time-integration rule.[13] The governing heat equation pairs the temporal rate of change of temperature to the second-order spatial derivatives (the Laplacian) and thus requires a computable approximation not only for $\frac{\partial u}{\partial n}$ but also for $\frac{\partial^2 u}{\partial x^2}$, $\frac{\partial^2 u}{\partial y^2}$ and $\frac{\partial^2 u}{\partial z^2}$. While forward Euler time stepping provides an approximation of the temporal derivative, it does not, by itself, deliver a mechanism to evaluate the spatial curvature terms that drive diffusion. By contrast, FDM replaces the spatial derivatives with grid-based finite-difference stencils, most commonly central differences, thereby converting the partial differential equation into an explicit algebraic update rule at each node. This discretization naturally leads to the FTCS scheme, which pairs forward-time stepping with central space differencing, resulting in a symmetric diffusion model consistent with isotropic (same temperature) heat spreading on a uniform mesh. However, it is important to note that the FTCS scheme employs a Taylor series expansion where higher-order terms are neglected. These neglected terms are the error of the large simulation that has compounded over long time ranges (over a large number of iterations), and thus this is an inherent uncertainty in this method.

1D baseline rod simulation:

Governing equation (1D Heat conduction with variable conductivity).

The one-dimensional temperature field $T(x,t)$, on a rod of length, L is governed by the transient heat conduction equation with spatially varying thermal conductivity $k(x)$:

$$\rho c_p \frac{\partial T}{\partial t} = \frac{\partial}{\partial x}(k(x)\frac{\partial T}{\partial x}) + q(x)$$

where ρ is density (kgm^{-3}), C_p is specific heat capacity ($Jkg^{-1}K^{-1}$), $k(x)$ is thermal conductivity ($Wm^{-1}K^{-1}$), and $q(x)$ is volumetric heat generation (Wm^{-3}). Because this paper only simulates the cooling of the brake discs, $q(x)$=0 such that the temperature decreases monotonically from the initial conduction under convective heat loss.

Geometry, grid, and material parameters:

A rod of length L=1.0 m was discretized using N_x=50 grid points:

$$\Delta x = \frac{L}{N_x - 1}, x_i = i\Delta x, i = 0, 1, \ldots, N_x - 1$$

Carbon-carbon composite representative constants used in the numerical model were:

$$\rho = 1780\ kgm^{-3}$$
$$c_p = 2218\ J\ kg^{-1}K^{-1}$$

Thermal conductivity was modeled as linearly varying along the rod

$k(x) = k_0(1 + 0.5\frac{x}{L}),\ k_0 = 86.3\ Wm^{-1}K^{-1}$

Such that conductivity increases from 86.3 to 129.45 $Wm^{-1}K^{-1}$ across the domain.

Boundary conditions (convective boundary conditions):

Convective heat exchange with ambient air was imposed at both ends of the 1D rod. With ambient temperature T_∞ and heat transfer coefficient h, the end conditions are:

At x=0, $k(0)\frac{\partial T}{\partial x}(0,t) = h(T(0,t) - T_\infty)$

At x=L, $-k(L)\frac{\partial T}{\partial x}(L,t) = h(T(L,t) - T_\infty)$

The simulations used T_∞=23 $°C$, and h=60 $Wm^{-2}K^{-1}$.

To represent a purely cooling scenario, the rod was initialized uniformly above ambient:

$$T(x,0) = T_0$$

with T_0 set at 1000 $°C$ (to allow for an analyzable difference between, which is higher than T_∞. With $q(x)$=0, this ensures the solution cools towards ambient over time rather than heating.

Numerical method: flux-conservative finite differences with FTCS time stepping:

The partial differential equation was solved numerically using the finite difference method on a uniform grid. A flux-conservative discretization was used for the variable-conductivity diffusion term by defining midpoint conductivities:

$$k_{i+0.5} = \frac{k_i + k_{i+1}}{2}$$

For interior nodes, i=1, ..., N_x-2 , the spatial operator was approximated as

Time integration was performed using an explicit Forward-Time Central-Space (FTCS) update:

$$T_i^{n+1} = T_i^n + \frac{\Delta t}{\rho c_p}[\frac{1}{\Delta x^2}(k_{i+0.5}(T_{i+1}^n - T_i^n) - k_{i-0.5}(T_i^n - T_{i-1}^n) + q_i]$$

For the cooling-only case, $q_i = 0$.

DOI: 10.36838/IJHSR820.18

Implementation of convective boundary conditions:

To enforce the convective boundary conditions within the explicit scheme, a ghost-node method was used. At x=0 , the derivative was approximated by a centered difference:[14]

$$\frac{\partial T}{\partial x}(0,t) \approx \frac{T_1 - T_{-1}}{2\Delta x}$$

Substituting into $k(0)T_x(0,t)=h(T_0-T_\infty)$ gives the left ghost value T_{-1} in terms of interior values:

$T_{-1} = T_1 - \frac{2\Delta x h}{k(0)}(T_0 - T_\infty)$

Similarly, at x=L,

$$\frac{\partial T}{\partial x}(L,t) \approx \frac{T_N - T_{N-2}}{2\Delta x}$$

And substitution into $k(L)T_x(L,t)=h(T_{N-1}-T_\infty)$, yields the right ghost value T_N. These ghost values were then used in the same flux-conservative formula to update the boundary nodes i=0 and i=N_x-1, ensuring the ends remain thermally active (the temperatures are not fixed but vary according to convection strength h).

Time stepping, stability control, and reported output times:

A uniform timestep Δt was used with explicit FTCS. Because explicit diffusion schemes are conditionally stable, timestep selection was controlled relative to grid spacing and the maximum effective diffusivity

$$\alpha_{\max} = \frac{max(k)}{\rho c_p}$$

and Δt was kept sufficiently small to satisfy the standard FTCS stability requirement, as $\Delta t \le \Delta t_{stable}$.[15]

Five temperature profiles were extracted and plotted at 30-second intervals such that the changes in temperature distribution would be more apparent and easier to identify.

All computations were implemented in Python (and ran in Google Colab), using array-based numerics and plotted as two-dimensional line graphs of $T(x,t)$ versus x. The rod domain was plotted over $0 \le x \le 1$ m for ease of simulation.

2D Carbon-carbon composite Plate cooling with convective (Robin) boundaries:

A two-dimensional transient heat-diffusion model was used to study cooling in a 0.5 x 0.5 m square-shaped carbon-carbon composite plate with convective heat loss at the outer perimeter and at the walls of the circular perforations within the boundaries of the plate. The plate temperature field $T(x,y,t)$ was governed by the 2D heat equation:

$$\frac{\partial T}{\partial t} = \alpha\left(\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2}\right)$$

where α is the thermal diffusivity.

Geometry and perforation configurations:

A square plate of side length 0.5 m was used for all 2D cases. Four configurations were simulated:

1) Unperforated plate, with a continuous solid domain with convection only at the edges of the plate

2) Singular central perforation of diameter 10 cm on the plate, with convection at the edges and perforation boundary

3) Perforated plate with 16 evenly distributed perforations of diameter 1 cm on the plate, with convection at the edges and the perforation boundary

4) Perforated plate with 25 evenly distributed perforations of diameter 1 cm on the plate, with convection at the edges and the perforation boundary

Boundary conditions:

Convective heat exchange was imposed using a convective (Robin) boundary condition:

$$-k\frac{\partial T}{\partial n} = h(T - T_\infty)$$

Applied on all outer edges of the plate and on internal boundaries along the perforation boundaries, where h is the convection coefficient, and n is the outward normal.

In the discrete scheme, convective boundary conditions were implemented using a Ghost-node method (where a virtual node outside the domain is used). To apply convective boundary conditions on a Cartesian grid without explicitly meshing curved boundaries, the boundary flux condition is enforced by replacing a missing neighbor (outside the domain) with an equivalent 'ghost' temperature. This is implemented in the form:[16]

$$T_{ghost} = T_{cell} - \beta(T_{cell} - T_\infty), \ \beta = \frac{h\Delta x}{k}$$

This is the discrete ghost-node expression consistent with enforcing the convective boundary condition at the boundary.

This is applied at the outer edges and the perforation boundaries in the simulations.

Numerical discretization: FTCS finite difference:

The 2D plate was discretized on a uniform Cartesian grid over [0, L_x] x [0, L_y] with Δx = Δy. The temperature was advanced using an explicit FTCS scheme:[17]

$$T_{ij}^{n+1} = T_{ij}^{n} + \alpha\Delta t\left(\frac{T_{i+1,j}^{n} - 2T_{ij}^{n} + T_{i-1,j}^{n}}{\Delta x^2} + \frac{T_{i+1,j}^{n} - 2T_{ij}^{n} + T_{i-1,j}^{n}}{\Delta y^2}\right)$$

with neighbor terms replaced by the convective ghost values whenever a boundary is encountered.

Stability condition:

The explicit diffusion scheme requires a time step satisfying the standard 2D criterion (r_x+ $r_y \le 0.5$); for Δx = Δy, this reduces to $r \le 0.25$ where $r = \frac{\alpha\Delta t}{\Delta x^2}$.[18]

This was enforced in the simulations using:

$\Delta t = 0.8\ \Delta t_{\max}$, $\Delta t_{\max} = 0.25\frac{\Delta x^2}{\alpha}$

Simulations:

Cells inside the perforations were treated as void and held at ambient (296.15 K):

$T(x,y,t) = T_\infty$ for perforations.

This provides a strong cooling sink and, combined with Robin handling at the solid-void interface, enforces internal convective heat removal at perforation walls.

High-resolution perforated cases (where circular perforations were modeled) were used N_x = N_y = 81 for the 16-perforation case for improved circularity of perforations on a Cartesian mesh.

Sampling of the thermal distribution of the four plates occurred at regular 30-second intervals, using a color scale for ease of analysis. 30 seconds was chosen, as it provided sufficient time for the temperature to change sufficiently such that a difference could be seen, but not too much time such that the temperature change would be too extreme.

To quantify the cooling rate, the solid-only mean temperature was computed at each time step:[19,20]

$$\underline{T}(t) = \frac{1}{A_{solid}} \int_{\Omega_{solid}}^{e} T dA \approx \frac{1}{N_{solid}} \sum_{(i,j)\in solid}^{3} T_{i\ j}$$

This was then plotted onto a graph for all cases (refer to Figures 2, 3, 4, and 5), allowing direct comparison.

All simulations and post-processing were implemented in Python using NumPy for array operations and Matplotlib for plotting/exporting figures (including multi-page PDF outputs). The scripts were executed in Google Colab to ensure reproducibility and consistent runtime settings.

Modeling assumptions

In this model, the material properties *a* and *k* are treated as constant, and heat transfer to the environment is represented solely by convective boundary conditions with constant *h* and T_∞, with radiation and contact resistance neglected. The fluid within the perforations is treated as an ambient reservoir at T_∞, while convective exchange at each perforation boundary is enforced on the solid side via the ghost-node method. Material properties within the plate are assumed homogeneous and isotropic, and no internal heat generation is included.

Parameters:

In this section, I list the parameters used for the 2D simulation. These values all reflect the properties of the carbon-carbon composite material that motorsport brake discs are composed of.

1) 2D Plate dimensions 0.5 x 0.5 m
2) Carbon-carbon composite density,[21] ρ = 1780.0 kgm^{-3}
3) Specific heat capacity,[22] C_p = 2218 $Jkg^{-1}K^{-1}$
4) Thermal conductivity,[23] k = 86.3 $Wm^{-1}K^{-1}$
5) Convective heat transfer coefficient, h = 100 $Wm^{-2}K^{-1}$
6) Initial plate temperature, T_{init} = 1273.15 K
7) Surrounding ambient temperature, T_∞ = 296.15 K

3D Transient brake disc cooling model:

The brake disc was modeled as a 3-dimensional annular solid using cylindrical coordinates, using the standard coordinate relations $x=r\cos\theta$, $y=r\sin\theta$, $z=z$, and discretized on a structured (r, θ, z) to support an explicit FTCS finite-difference time integration. Heat transfer within the disc was governed by the transient heat equation

$\frac{\partial T}{\partial t} = \alpha(T)\nabla^2 T$, where $\alpha(T) = \frac{k(T)}{\rho c_p(T)}$

The specific heat capacity, $C_p(T)$ was implemented as a piecewise-linear interpolation table, using the three tabulated points at 473.15 K, 1273.15 K, 1773.15 K. The thermal conductivity, $k(T)$ was also treated as temperature-dependent using a piecewise-linear table, and the thermal diffusivity was updated consistently as $\alpha(T) = \frac{k(T)}{\rho c_p(T)}$.

Geometry of the brake disc:

The disc geometry was defined by the outer diameter D_O= 0.278 m, inner diameter D_i = 0.1816 m, and thickness H= 0.028 m, giving $R_O = D_O/2$ = 0.139 m and $R_i = D_i/2$ = 0.0908 m. Because the inner radius of the brake disc is not specified in the 2014 FIA Regulations,[24] and F1 teams prefer to keep their designs secret to maintain a competitive edge, D_i , and thus R_i , was estimated from an image of the Mercedes W05 at the German Grand Prix.[25] Through-radial perforations were represented with an estimated perforation diameter d_p=0.003 m and perforation depth set equal to the full radial width of the annulus, $L_{depth} = R_O - R_i$.[26]

Material model:

The disc is treated as an effective isotropic carbon-carbon composite solid with constant density ρ = 1780 kgm^{-3} and temperature-dependent heat capacity and conductivity. The specific heat capacity, $C_p(T)$ was implemented as a piecewise-linear interpolation table, using the three tabulated points at 473.15 K, 1273.15 K, 1773.15 K.

The material properties were updated one per global time-step using the volume-weighted mean temperature $T_{mean}(t^n)$ to maintain internal consistency in the transient property model $c_{p,eff}{}^n = c_p(T^n{}_{mean})$, $k^n{}_{eff} = k(T^n{}_{mean})$, $\alpha^n{}_{eff} = \frac{k^n{}_{eff}}{\rho\, c^n{}_{p,\,eff}}$.

Initial conditions:

The disc is initialized at a uniform temperature representing the post-braking thermal state of the brake disc of a 2014 Mercedes W05 F1 car exiting turn 2 at the Autodromo Nazionale Monza Circuit in Italy, where $T(r, \theta, z, \theta) = T_0$ = 1000 $^\circ C$ = 1273.15 K. Ambient air temperature is prescribed as T_∞ = 23 $^\circ C$ = 296.15 K.

This uniform initial temperature assumption provides a controlled starting condition, allowing the subsequent cooling to be attributed primarily to conduction, convection, and perforation-enhanced heat loss during the 15 second acceleration window.

Boundary conditions:

Cooling to airflow is represented using convective boundary conditions, which model heat loss proportional to the temperature difference between the solid and the surrounding air at convective boundaries. Accordingly, convective boundary conditions were applied on the inner bore, $r=R_i$; Outer rim, $r=R_O$ and both disc faces, z=0 and $z=H$, in the standard form:

For numerical implementation within the explicit FTCS scheme,[29] the convective boundary condition

$$-k_{eff}\frac{\partial T}{\partial n} = h(t)(T - T_\infty)$$

is enforced using a ghost-cell formulation, where n is the outward normal, T_∞ is ambient temperature, and $k_{eff} = k(T_{mean})$ is the effective conductivity evaluated at the volume-weighted mean temperature at each time step. This ghost-cell approach converts the boundary heat-flux relation into an equivalent virtual temperature just beyond the computational domain, enabling the same centered-difference Laplacian stencil to be applied at boundary nodes whilst still imposing convective heat loss.

DOI: 10.36838/IJHSR820.18

Radial boundaries ($r=R_i$, $r=R_O$):

At the inner bore, $r=R_i$, the outward normal points in the $-\hat{e}_r$ direction, so $\frac{\partial T}{\partial n} = -\frac{\partial T}{\partial r}$. The convective boundary condition becomes

$$-k_{eff}\frac{\partial T}{\partial n} = h_{in}(t)(T - T_\infty)$$

Using a centered difference at the boundary node with ghost temperature T_{-1} gives

$$\frac{T_1 - T_{-1}}{2\Delta r} = \frac{h_{in}(t)}{k_{eff}}(T_0 - T_\infty) \;\rightarrow\; T_{-1} = T_1 - 2\Delta r\frac{h_{in}(t)}{k_{eff}}(T_0 - T_\infty)$$

At the outer rim $r=R_O$, the outward normal is $+\hat{e}_r$, thus $\frac{\partial T}{\partial n} = \frac{\partial T}{\partial r}$ and

$$-k_{eff}\frac{\partial T}{\partial n} = h_{out}(t)(T - T_\infty)$$

With boundary node T_{N_r-1} and ghost node T_{N_r}

$$\frac{T_{N_r} - T_{N_r-2}}{2\Delta r} = -\frac{h_{out}(t)}{k_{eff}}(T_{N_r-1} - T_\infty) \longrightarrow T_{N_r} = T_{N_r-2} - 2\Delta r\frac{h_{out}(t)}{k_{eff}}(T_{N_r-1} - T_\infty).$$ [30]

These ghost values are substituted into the FTCS Laplacian evaluation such that the radial diffusion term remains second-order in space whilst consistently applying the convective flux at both radial edges.

Axial (face) boundaries:

Similarly, convection on the disc faces is enforced using ghost temperatures with outward normal $-\hat{e}_z$ at $z=0$ and $\hat{e}_z$ at $z=H$. The ghost values used in the centered difference are

$$T_{z,-1} = T_{z,1} - 2\Delta z\frac{h(t)}{k_{eff}}(T_{z,0} - T_\infty),\; T_{z,N_z} = T_{z,N_z-2} - 2\Delta z\frac{h(t)}{k_{eff}}(T_{z,N_z-1} - T_\infty).$$ [31]

Consistency with the time-dependent convection model:

At each time step, $h(t)$, $h_{in}(t)$ and $h_{out}(t)$ are evaluated from the prescribed velocity-based convection model, and k_{eff} is obtained from the temperature-dependent material model. This ensures that the applied boundary heat flux remains consistent with both the time-varying airflow during the acceleration of the F1 car and the temperature dependence of the carbon-carbon composite properties, specifically thermal conductivity, specific heat capacity, and thermal diffusivity.

Ghost-cell temperature is recomputed each time step using the current $h(t)$ and k_{eff} and substituted into the boundary Laplacian stencil.

Perforation-enhanced cooling modeling:

Because the brake discs of 2014 F1 cars contain perforations (Ventilation holes), additional cooling is included through 2 physically motivated mechanisms without explicitly meshing each perforation.

1) The internal perforation-wall convection as a distributed sink term

The increased internal surface area provided by the perforation walls was modeled as a volume-averaged (represented as an effective volumetric cooling term) convective heat-loss term applied throughout the solid:[28]

$$\frac{\partial T}{\partial t} = -\beta(t)(T - T_\infty),\; \beta(t) = \frac{h(t)\,A_{perforations}}{\rho\, c_p\, (T_{mean})V_{solid}}$$

where $\beta(t)$ is the time-dependent volumetric cooling rate coefficient, $A_{perforations}$ is the total lateral area of all cooling perforations (where the endcaps are neglected). ($A_{perforations} = N_{perforations}\pi d_h L_{depth}$) and V_{solid} is the net solid material volume of the disc ($V_{solid} = V_{annulus} - V_{perforations}$).

$$V_{annulus} = \pi(R_o{}^2 - R_i{}^2)H,\; V_{perforations} = N_{perforations}\pi(\tfrac{d_h}{2})^2 L_{depth}$$

This sink term (added volumetric cooling term) is integrated analytically over each Strang half step $\Delta t/2$

$$T \leftarrow T_\infty + (T - T_\infty)e^{-\beta(t^n)\frac{\Delta t}{2}}$$

2) Perforation-induced enhancement of convection at the radial-edge:

For perforations (through the radius), the perforations also increase airflow interaction near both radial edges.[1] This is represented by boosting the effective convection coefficient on the inner bore and outer rim by distributing the perforation (hole) area evenly between the edges:

$$h_{edge}(t) = h(t)\frac{A_{edge} + 0.5A_{perforations}}{A_{edge}},\; A_{edge}(R) = 2\pi RH$$

Applied separately at $R=R_i$ and $R=R_O$. The face convection coefficient $h(t)$ used on $z=0$ and $z=H$ was not edge boosted. A modeling approximation, the factor $0.5A_{perforations}$ assumes that the through-radial perforations contribute equally to enhanced convection at the inner and outer radial edges. (The face coefficient $h(t)$ applied on $z=0$ and $z=H$ was not edge boosted, matching the implementation).

Together, the convective boundary conditions and the perforations allow the model to represent both external forced convection and perforation-enhanced heat removal.

Airflow convection-coefficient model:

To represent increasing forced convection during acceleration out of turn 2 in Monza, the convection coefficient is modeled as a function of time via a velocity series, prescribing a piecewise-linear velocity profile, defined at 1 second intervals over 0-15 *s*. The prescribed velocity (V/kmh^{-1}) values over 0-15 *s* is shown in the table below.[32]

Table 1: Mercedes W05 2014 F1 car Speed-Time relationship out of turn 2 Monza.

t/s	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
v	67	80	138	192	225	250	268	283	290	300	310	315	325	328	332	332

These values were used as a representative for the magnitude of forced convection.

A power law correlation maps the changing velocity to convection:[33]

$$h(t) = h_{ref}(\frac{v(t)}{v_{ref}})^{0.8}$$

with $h_{ref} = 100 Wm^{-2}K^{-1}$ at $V_{ref} = v(0)$. This provides a simple monotonic increase in convection strength during acceleration, consistent with the modeling approach of forced-air cooling.

Numerical method:

The computational domain for the brake disc was defined in cylindrical coordinates (r, θ, z) over the annular region $R_i \le r \le R_O$, $0 \le \theta \le 2\pi$, and $0 \le z \le H$. The temperature field $T(r, \theta, z, t)$ was discretized on a structured cylindrical grid with resolutions $N_r = 220$, $N_\theta = 360$, $N_z = 64$ such that $T_{i,j,k}{}^n \approx T(r_i, \theta_j, z_k, t_n)$. The

nodal coordinates were generated as $r_i = R_i + i\Delta r$, $i=0,\ldots,N_r-1$, $\theta_j = j\Delta\theta$, $j=0,\ldots,N_\theta-1$, $z_k = k\Delta z$, $k=0,\ldots,N_z-1$, with uniform spacings $\Delta r = \frac{R_o - R_i}{N_r - 1}$, $\Delta\theta = \frac{2\pi}{N_\theta}$, $\Delta z = \frac{H}{N_z - 1}$.

Periodicity was imposed in the circumferential direction to represent the full annulus:

$$T(r, 0, z, t) = T(r, 2\pi, z, t)$$

which is implemented numerically by wrap-around neighbors in the θ-index during the centered-difference evaluation of $\frac{\partial^2 T}{\partial \theta^2}$.[34]

The 3-dimensional transient temperature field $T(r, \theta, z, t)$ was advanced using an explicit Forward-Time Centered-Space (FTCS) finite-difference scheme applied to the cylindrical Laplacian.[35] The explicit update at the interior nodes is written as

$$T_{i,j,k}{}^{n+1} = T^n{}_{i,j,k} + \alpha^n{}_{eff}\Delta t(\nabla^2 T)^n{}_{i,j,k}$$

where $(\nabla^2 T)^n{}_{i,j,k}$ is evaluated using second-order centered differences in r, θ, and z,[36] consistent with the cylindrical Laplacian, and $\alpha^n{}_{eff} = \frac{k^n{}_{eff}}{\rho c^n{}_{p,eff}}$ is the effective thermal diffusivity evaluated at the representative disc temperature $T_{mean}(t_n)$ where $k^n{}_{eff} = k(T_{mean})$ and $c^n{}_{p,eff} = c_p(T_{mean})$.[37] Forced convection to air was imposed through convective boundary conditions on the inner bore ($r=R_i$) , outer rim ($r=R_O$), and both disc faces ($z=0$, $z=H$). Within the explicit FTCS Formulation, these convective conditions were implemented using a ghost-cell formulation, which defines virtual temperatures just outside the domain such that the centered-difference Laplacian stencil can be applied at boundary nodes while enforcing the convective heat flux.

Perforation-enhanced cooling was incorporated without explicitly meshing individual perforations by coupling diffusion with an internal convective sink term via Strang Splitting. At each global time step, Δt, the update follows the symmetric operator-split sequence:[38]

1) Internal sink half-step over $\Delta t/2$ applied analytically as $T < -T_\infty + (T - T_\infty)e^{(-\frac{\beta(t)\Delta t}{2})}$ where $\beta(t) = \frac{h(t)\, A_{perforations}}{\rho\, c^n{}_{p,eff} V_{solid}}$.

2) Diffusion full-step over Δt using FTCS with centered spatial differences in (r, θ, z,) and periodicity in θ.

3) Internal sink half-step over $\Delta t/2$, again applied using the same exponential update.

To ensure numerical stability of the explicit diffusion update at the chosen spatial resolution, the diffusion full-step is performed with internal sub-cycling: each global step Δt is subdivided into n_{sub} smaller diffusion steps $\Delta t_{sub} = \Delta t / n_{sub}$ that satisfies a conservative FTCS Stability bound.

At each time step, the temperature-dependent properties are updated consistently by evaluating $c_p(T)$ and $k(T)$ from the prescribed piecewise-linear tables at the volume-weighted mean disc temperature $T_{mean}(t^n)$, then computing $c^n{}_{p,eff} = c_p(T_{mean})$, $k^n{}_{eff} = k(T_{mean})$, and $\alpha^n{}_{eff} = \frac{k^n{}_{eff}}{\rho c^n{}_{p,eff}}$.[39] The time-varying convection coefficient $h(t)$ is obtained from the prescribed velocity series via the power-law model, and effective edge coefficients $h_{in}(t)$ and $h_{out}(t)$ are computed when perforation edge-enhancement is applied.

Outputs:

For each simulation case, the solver produces a single multi-page PDF containing absolute mid-plane temperature contours at 1 s intervals from $t=0$ to $t=15$ seconds. The mid-plane is defined at $z=H/2$, because $N_z=64$ is even, the true mid-thickness plane lies between the 2 central z-slices, thus the plotted mid-plane field is taken to be the average of the 2 central slices. Each of the 2 PDF outputs is a plane-view of the annulus (a carbon-carbon composite brake disc with 800 and 1000 perforation jacket) on the x-y plane, with the inner and outer radii overlaid, and includes a colorbar in Kelvin. The page title reports time t, the predetermined velocity $v(t)$, the base face convection coefficient $h(t)$, and the case identifier.

At each global time step Δt, the update proceeds as an Internal sink half-step over $\Delta 0.5t$, diffusion full-step over.

Results and Discussion

1D Baseline: Variable-conductivity rod with internal heat source:

A one-dimensional rod model was implemented as a baseline case to isolate conductive heat redistribution and end-losses under convection. The domain was a rod of length $L=1$ m, discretized using $N_x=50$ uniformly spaced nodes. The material properties were taken as representative carbon-carbon composite values and implemented in the transient heat equation with convection at both ends:

$$k(0)\frac{\partial T}{\partial x}(0,t) = h(T(0,t) - T_\infty),\ -k(L)\frac{\partial T}{\partial x}(L,t) = h(T(L,t) - T_\infty)$$

where T_∞ is the ambient temperature, and h is the heat-transfer coefficient. Numerically, the partial differential equation was solved using an explicit FTCS finite-difference scheme, with a flux-conservative spatial discretization and convective boundary conditions enforced using a ghost-node formulation such that the boundary heat flux responded to (T-T_∞) rather than fixing the boundary temperature.

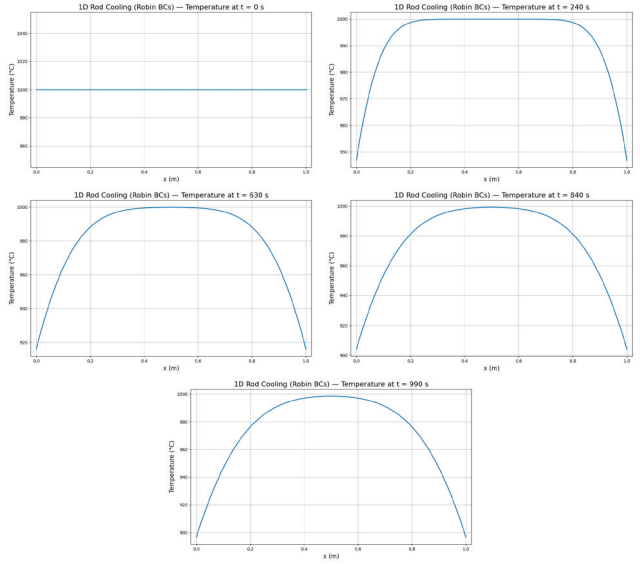


Figure 1: 1D carbon-carbon composite rod temperature profile T(x,t) under convective boundary conditions at both ends, showing the spatial distribution after the reported simulation time. The rod cools towards the ambient temperature through boundary heat exchange governed by h.

DOI: 10.36838/IJHSR820.18

Figure 1 shows a cooling profile consistent with a symmetric setup (uniform initial temperature of 1000 $^{\circ}C$ and identical convection at both ends). The temperature remains the greatest in the center while the ends cool more rapidly due to the convective boundaries permitting the loss of heat from the domain in proportion to (T-T_{∞}). As time increases, the thermal gradient near each boundary steepens, then gradually decreases as the entire rod approaches T_{∞}(23 $^{\circ}C$). This baseline case ensures that the numerical discretization and Robin boundary implementation produce physically expected behavior.

2D temperature-field evolution under convective (Robin) cooling:

A carbon-carbon composite square plate of side length 0.5 *m* was simulated using the transient two-dimensional heat equation with convective (Robin) boundary conditions applied to exposed boundaries. Across all configurations, temperature fields were depicted at regular time intervals, simulating the decrease in the temperature of the plate. The investigated geometries were: (1) a plate with no perforations, (2) a plate with a singular, central 10 *cm* diameter perforation at the center of the plate, (3) a plate with 25 regularly arranged, evenly divided perforations with a diameter of 1 *cm* , (4) a plate with 16 regularly arranged, evenly divided perforations with a diameter of 1 *cm*.

Solid carbon-carbon composite plate (with the outer edges having convective boundary conditions):

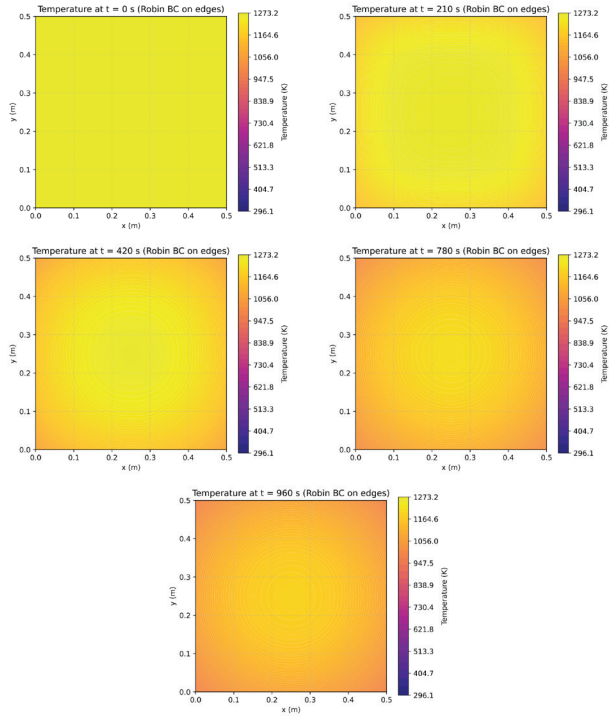


Figure 2: Solid carbon-carbon composite plate simulation with temperature maps at 30 *s* intervals from 0-960 *s* with convective cooling on all outer edges.

This baseline result, depicted in Figure 2, shows the expected diffusion-dominated behavior: the strongest gradients develop at the external boundaries where convective heat loss occurs, while the interior remains hotter for longer because heat must be conducted to the perimeter before it can be removed. As time increases, the temperature field becomes progressively smoother as diffusion reduces gradients throughout the domain, and the solution approaches a low-gradient state consistent with sustained convective cooling at the edges.

Carbon-carbon composite plate (with a singular central 10 cm diameter circular perforation at the center of the plate)

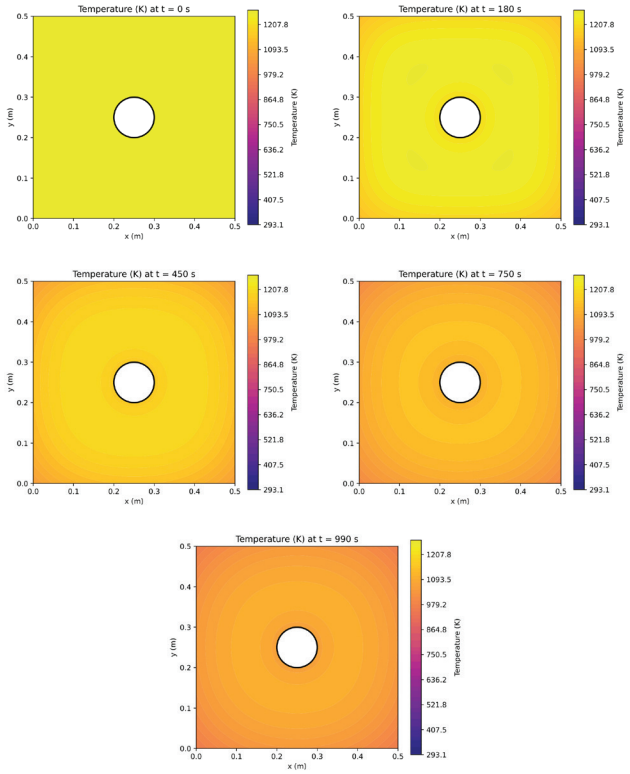


Figure 3: Single-holed (single perforation) carbon-carbon composite plate simulation (with the outer edges and the inner perforation-fluid boundary having convective boundary conditions).

Introducing a single circular perforation at the center of the plate adds an internal solid-fluid interface where convective boundary conditions are also applied (along with those at the outer edges of the plate). As such, the isotherms (lines on a graph indicating equal temperature) are not solely controlled by distance to the external perimeter; rather, the interior temperature structure is shaped by simultaneous conduction inwards towards the central perforation and conduction outwards towards the outer edges of the plate, as shown in Figure 3. However, as time increases, diffusion decreases, and the field becomes smoother.

Carbon-carbon composite plate (with 16 circular perforations with a 1 cm diameter arranged regularly):

This configuration reduces the characteristic conduction distance from interior points to a convective boundary, which promotes a more spatially distributed cooling pattern. Over long times, diffusion still smooths the field, but the approach to a low-gradient state is mediated by several heat-exchange locations evenly distributed across the plate, as depicted in Figure 4.

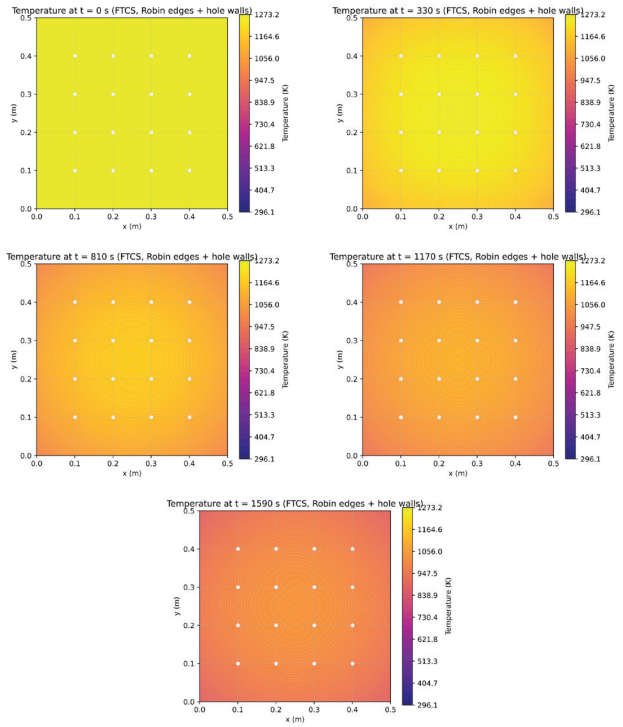


Figure 4: 16-holed(16-perforation) carbon-carbon composite plate simulation (with the outer edges and the 16 inner perforation-fluid boundaries having convective boundary conditions).

Carbon-carbon composite plate (with 25 circular perforations with a 1 cm diameter arranged regularly)

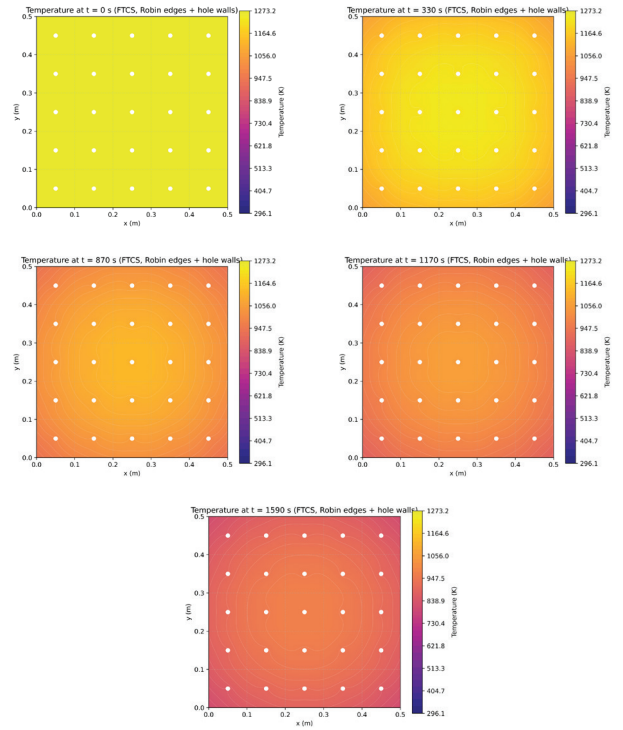


Figure 5: 25-holed (25-perforation) carbon-carbon composite plate simulation (with the outer edges and the 25 inner perforation-fluid boundaries having convective boundary conditions).

Compared to the plate with 16 perforations, a more even pattern of heat distribution as many local convective interfaces influence the temperature simultaneously. However, the presence of many (25) evenly distributed sinks also means that large interior regions are no longer insulated from convection by long conduction paths to the perimeter, as evidenced by Figure 5. Over time, diffusion homogenizes the field, but the trajectory toward a low-gradient, near-ambient state is driven by heat removal at numerous internal boundaries rather than predominantly at the external edges.

Physical application of this simulation:

Within this simplified model, the simulations demonstrate that perforations within the material can act as distributed thermal sinks, reshaping transient temperature patterns. This supports the design intuition that perforations do more than reduce mass: they also alter the thermal pathway topology by adding internal convective boundaries where heat loss can occur.

A 25-perforation configuration cools faster than a 16-perforation configuration as it introduces more thermally active internal boundaries and also reduces the distance from any point on the plate to a boundary (assuming the perforations are evenly distributed). This is evident in Figure 6 (below) when the orange line indicating 25 perforations has a smaller gradient than the blue line indicating 16 perforations, therefore showing that the greater the number of perforations, the higher the rate of cooling.

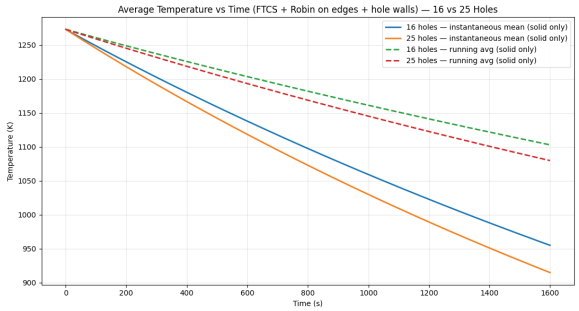


Figure 6: Combination of the graphs of Temperature against time for 16 and 25 (holes) perforations (Note: the x-axis is very elongated, thus it may be difficult to view the graph as nonlinear).

3D FTCS Brake disc cooling:

Two perforation cases were simulated for the Mercedes W05 brake disc under the same 15 *s* acceleration-cooling period, where the car reaches maximum speed at 14 *s*: 800 perforations and 1000 perforations. In both simulations, the disc was initialized at a uniform temperature T_0 and cooled toward ambient T_∞ through forced convection applied on the inner bore, outer rim, and both brake disc faces, with perforation-enhanced cooling represented by an internal volumetric sink and an edge convection boost.

Figure 7 and Figure 8 present the peak and volume-mean temperature histories predicted by the 3-dimensional FTCS Carbon-carbon composite brake disc model over 15 *s* during the simulated acceleration of the Mercedes W05 out of turn 2 in the Autodromo Nazionale Monza Circuit. For both the 800-perforation and 1000-perforation simulations, the peak temperature remains above the volume-mean temperature throughout the simulation, indicating that spatial non-uni-

DOI: 10.36838/IJHSR820.18

formity persists as the disc cools. The non-linear curvature of the graphs is consistent with the imposed time-dependent forced-convection: As the prescribed airflow increases with time, the effective convective heat removal increases, accelerating the decay of both the peak and mean temperatures. This curvature of the 2 simulations reflects competing effects. Because $h(t)$ varies in time, while the driving temperature difference (T-T_∞) decreases as the disc cools, weakening the instantaneous convective driving, with further modulation from $C_p(t)$ and $k(t)$.

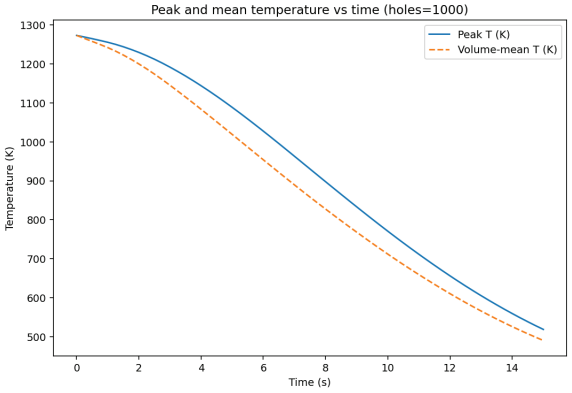


Figure 7: Peak and volume-mean disc temperature against time graph for the 1000-perforation (hole) annular disc case over 15 *s* under velocity-driven convection, in Kelvin.

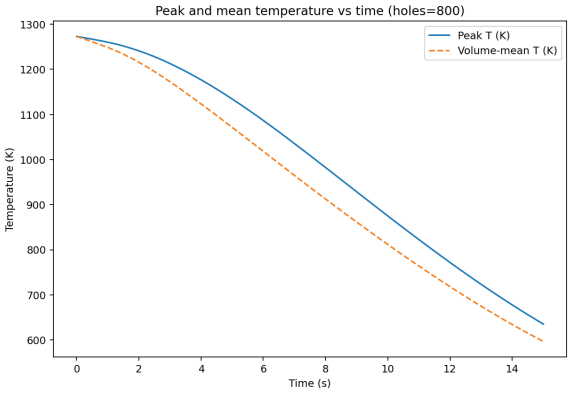


Figure 8: Peak and volume-mean disc temperature against time graph for the 800-perforation (hole) annular disc case over 15 *s* under velocity-driven convection, in Kelvin.

A clear geometric effect is observed when comparing the 2 graphs. The 1000-perforation case cools more rapidly than the 800-perforation case at all times t , with the difference in the temperature increasing near the end of the simulations as the cumulative effect of the stronger internal sink (of the 1000-perforation simulation) increases. By $t \approx 15\ s$, the 1000-perforation simulation reaches approximately $T_{peak} \approx 5.2 \times 10^2 K$, whereas the 800-perforation case remains higher at approximately $T_{peak} \approx 6.3 \times 10^2 K$, a difference of approximately 100 K when comparing variables for each simulation. This difference arises from the increased effective internal convective area and associated volumetric sink strength in the higher perforation disc, due to the increased heat-transfer surface area, which increases the net heat loss from the solid brake disc without explicitly meshing individual perforations.

The separation between peak and mean temperature also provides a compact indicator of thermal gradients: where the gap is greater, the field is more heterogeneous, suggesting a greater likelihood of local hot-spot persistence even as the bulk temperature decreases.

From a practical standpoint, under identical airflow forcing and material property updates, the higher the number of perforations (and thus the higher the total surface area) of the brake disc, the higher the overall rate of cooling of the brake disc, decreasing the peak temperature over the 15 *s* acceleration (and thus cooling) time period, suppressing the hottest regions more effectively over this period of time, suggesting a greater thermal-management margin for the more heavily perforated brake disc. This is evident in Figure 9 (below).

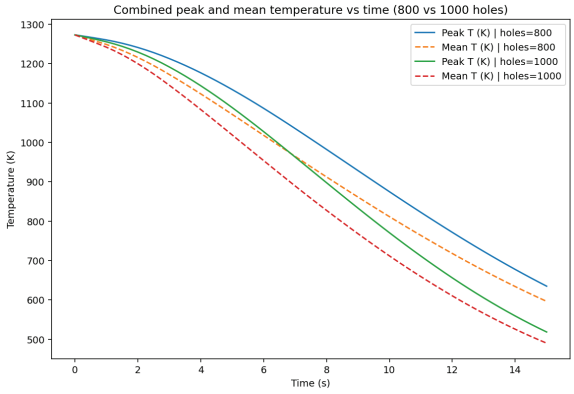


Figure 9: Combination of the graphs of Temperature against time for the 800-perforation and 1000-perforation simulations.

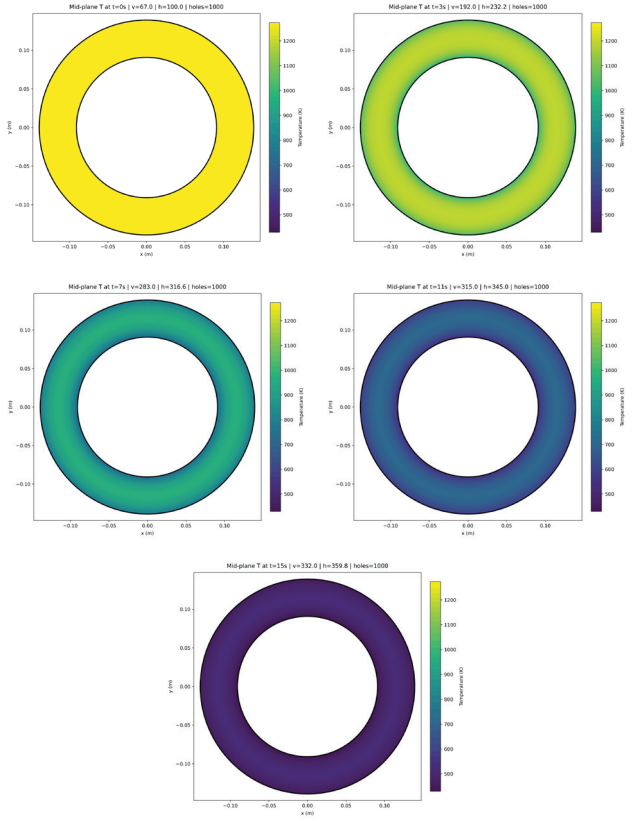


Figure 10: Mid-plane temperature contours for the 1000-perforation annular brake disc over 15 *s* under velocity-driven convection.

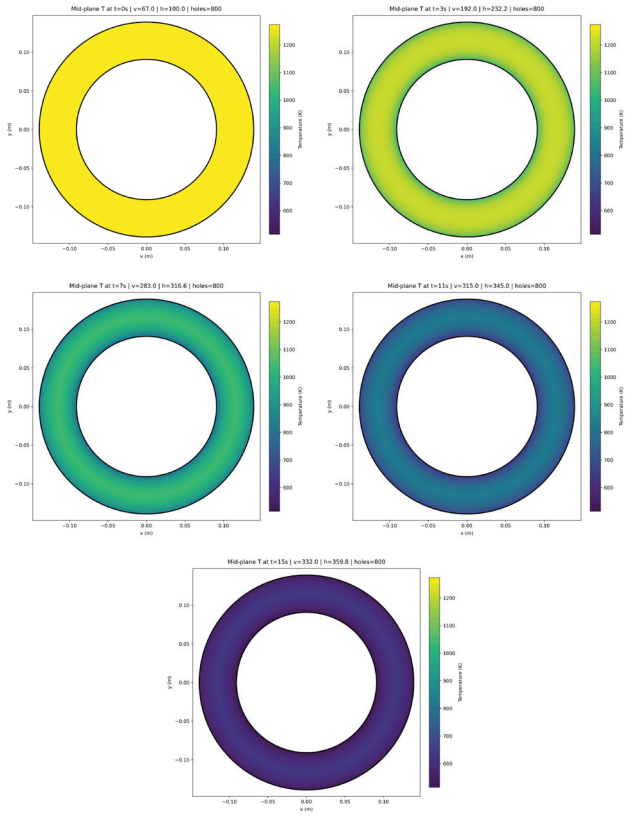


Figure 11: Mid-plane temperature contours for the 800-perforation annular brake disc over 15 *s* under velocity-driven convection.

Across both perforation configurations, the multi-page mid-plane temperature field outputs (figures 4 and 5) indicate the qualitative cooling topology. At time t=0 s, both cases are spatially uniform at the prescribed initial temperature, T_0 = 1273.15 K, with cooling first developing as radial boundary layers adjacent to the inner bore and outer rim before progressively penetrating toward mid-radius as time increases. This is consistent with the imposed convective boundary conditions at the edges of the brake disc and faces of the brake disc, and with the monotonic increase in the prescribed airflow-driven convection during the acceleration of the W05 F1 car out of turn 2 in Monza. Simultaneously (as time advances), the fields remain circumferentially uniform, which is expected because the boundary conditions are applied uniformly around the brake disc and periodicity is enforced in θ. The dominant spatial structure is thus radial, with the inner bore and outer rim cooling the fastest, producing cooler bands at r=R_i and r=R_O, while the interior retains higher temperature at early times. This behavior is consistent with convection acting on all exposed boundaries, with an increasing forced convection intensity through the prescribed $h(t)$ during the 0-15 *s* acceleration period.

When the 2 simulations are compared directly, the 1000-perforation configuration cools more rapidly than the 800-perforation configuration under the same external airflow, at the same time. The physical reason is geometric: increasing the number of through-radial perforations increases the total internal convective surface area used by the distributed sink term and also strengthens the effective edge convection through the edge boost allocation, so the disc experiences a larger net heat removal capacity for the same external airflow history. This difference is also visible in the mid-plane contours, where the 1000-perforation simulation exhibits an earlier and more extensive shift towards cooler color levels at the same time compared to the 800-perforation simulation. Interpreted physically within the model assumptions, increasing perforation count increases the number of thermally active internal boundaries, via the sink term, and enhances edge convection, reducing both the absolute temperature level and the persistence of hot regions during rapid post-braking cooling.

Model limitations:

1) Thermal radiation is not modeled in the simulation:

Adding a realistic radiation model to my simulation involves not just introducing T^4 on the surface. In practice, radiation can require nonlinear boundary terms, temperature-dependent emissivity, and coupling with convection and conduction at every step. With the already demanding FTCS 3D grid, adding radiative exchange would substantially increase the runtime and memory usage and therefore would require much more powerful computing devices than what is currently available. In fact, major companies such as AWS and Oracle partner with F1 teams, providing services for computational analysis of their Formula 1 cars.

Overall, however, with the current high-resolution 3D explicit FTCS framework, incorporating coupled enclosure radiation would substantially increase computational cost and implementation complexity, often motivating the use of high-performance computing resources for iterative design-grade studies.

Additionally, due to the compact nature of the modern drum-shaped F1 brake system, any heat radiated from the brake disc will be absorbed by this dark colored carbon-fiber drum and will be re-radiated back towards the brake disc. As such, there should not be much variation between the current simulation employed and the simulation updated with a radiation model.

For future work, one may consider non-linearities such as the Stefan-Boltzmann Law, which accounts for thermal radiation.

2) The convection coefficient h(t) is a simplified surrogate:

The convection coefficient is prescribed through a power-law dependence on a velocity trace. This captures the trend of increasing forced convection during acceleration, but it does not fully resolve full under-wheel aerodynamics, brake duct geometry, or turbulence structure.

This simplified convection model is employed because realistic forced convection around an F1 brake disc is strongly governed by 3-dimensional unsteady flow: Rotating wheel pumping, wheel-wake interaction, brake-disc jet impingement, strong curvature and separation, and turbulence. In a higher-resolution, higher-fidelity approach, the local heat transfer coefficient h is not prescribed; it emerges from a conjugate heat transfer computation where the airflow and disc conduction are solved in a coupled manner, producing a spatially and temporally varying all heat flux. Achieving this fidelity typically demands fine near-wall meshing, turbulence modeling, and

DOI: 10.36838/IJHSR820.18

small time steps, leading to large cell counts and long runtimes, a simulation that would require the use of powerful computational devices.

Therefore, h is prescribed in my model as an engineering surrogate to capture the first-order effect of increasing forced convection during acceleration while keeping the model computationally tractable.

3) Perforations in the brake disc are not explicitly meshed:

Approximating the perforations as a distributed sink and an edge boost avoids meshing thousands of features, yet it fails to capture local jetting and recirculation.

The unmeshed model in the simulation is used, since meshing thousands of perforations substantially increases the grid size and runtime. To resolve a perforation with diameter Several cells across it in each direction are necessary. Repeating this process for up to 1000 perforations increases the grid into hundreds of millions of cells, which is unfeasible without the use of powerful computing devices.

4) Effective isotropic conductivity and representative temperature:

In this model, the temperature-dependent thermal conductivity $k(T)$ and specific heat $c_p(T)$ are implemented as piecewise-linear tables, but (for computational tractability in the High-resolution 3-dimensional explicit FTCS scheme) they are evaluated once per time step using the volume-weighted mean temperature $T_{mean}(t)$. This gives stepwise-constant effective values $k_{eff} = k(T_{mean})$ and $c_{p,eff} = c_p(T_{mean})$ that are applied uniformly across the domain for that time step. However, this model forces the whole disc to use the same same k_{eff}, and $c_{p,eff}$ based on T_{mean}, so hotspots and cooler regions cannot conduct or store heat differently, despite $k(T)$ and $c_p(T)$ are temperature-dependent. Consequently, the model can misestimate local cooling rates and temperature gradients, reducing accuracy for predicting hotspot persistence and subtle small-scale thermal non-uniformity.

The model employs this simplified material behavior because updating $k(T)$ and $c_p(T)$ everywhere would be computationally expensive. The simulation evaluates properties at the volume-weighted mean temperature, once per global step, and uses that for k_{eff}, $c_{p,eff}$ and α_{eff} . This prevents the necessity of recalculating variable coefficients at every node sub-step, which would significantly raise the computational cost whilst still maintaining a relatively high degree of accuracy.

Physical implications of this simulation:

Within this simplified FTCS-based 3-dimensional model, a comparison between the 1000-perforation and 800-perforation results from this simulation indicates a consistent acceleration of the rate of cooling when the number of perforations is increased, evidenced both by the peak temperature and volume-mean temperature plots in Figure 7 and Figure 8. In both simulations, the curves are monotonically decreasing over the 0-15 second window, but the 1000-perforation case remains systematically lower after the initial transient, indicating that the additional perforation area increases the resultant heat-removal rate through the combined mechanisms used in the model.

By the end of the simulated interval, the 1000-perforation case reaches a peak temperature of approximately 520 K and a volume-mean temperature of approximately 500 K , whereas the 800-perforation case remains higher at roughly 630 K and and a volume-mean temperature of 600 K. This corresponds to an end-of-window reduction on the order of $\Delta T \approx 110$ K in both peak and mean when moving from 800 to 1000 perforations.

In a 2014 Formula 1 operating context, that scale of reduction is sufficiently large to be meaningful: it implies less time spent at elevated brake disc temperatures during the corner-exit acceleration phase, which decreases the probability of brake overheating and thus brake failure, thus increasing repeated braking performance, and reduces the possibility of brake overheating (despite the high maximum temperature which can be withstood by carbon-carbon composite brakes) and brake failure.

■ Conclusion

This study developed a progressively higher-fidelity transient heat-conduction framework for motorsport brake cooling. Beginning with 1-Dimensional and 2-Dimensional FTCS finite-difference simulations with convective (Robin) boundary conditions to confirm that the explicit FTCS diffusion scheme behaves correctly via the 1D case. It also ensured that the scheme is applicable to a larger field with geometric effects, before simulating the computationally expensive 3-Dimensional simulation. In the 3D simulation, the carbon-carbon composite disc representative of the 2014 Mercedes W05 Formula 1 racecar was initialized uniformly at T_0 = 1273.15 K and cooled over 15 s, towards ambient air T_∞ = 296.15 K using time-dependent convective boundary conditions on the inner bore, outer rim and both faces of the brake disc, where increasing convection coefficient $h(t)$ is driven by the prescribed, increasing velocity history; the mid-plane fields and peak volume-mean graphs depict boundary-initiated cooling with persistent internal non-uniformity over this period of time. Under identical material parameters and conditions, the 1000-perforation configuration cools at a higher rate and reaches lower peak and mean temperatures compared to the 800-perforation configuration, as the modeled effective convective area (the total surface area) of the brake disc is greater due to the presence of an additional 200 perforations for the 1000-perforation configuration simulation (compared to the 800-perforation configuration simulation). Therefore, according to the model employed, it follows that the greater the surface area in contact with the surrounding fluid, the greater the rate of cooling. In motorsport, namely, Endurance racing, Indycar and Formula 1, where brake thermal margins are narrow, as engineers are persistently working to shave hundredths of a second off of lap times, and because consistent high performance braking is repetitive, these simulations are significant as improved heat loss rate can stabilize brake response when

there are frequent heavy braking zones and reduce the risk of thermally driven performance loss across each lap.

Acknowledgments

I would like to thank my parents for their constant support throughout this research project. I also sincerely thank Dr Adam Almakroudi, Researcher at Imperial College London, for mentoring me throughout the research process and elevating my understanding of the topic at hand.

References

1. Belhocine, A.; Abu Bakar, A. R.; Bouchetara, M. Numerical Modeling of Disc Brake System in Frictional Contact. *Tribology in Industry* **2014**, *36* (1), 49–66. https://www.tribology.rs/journals/2014/2014-1/7.pdf.
2. Bellini, C.; Di Cocco, V.; Iacoviello, D.; Iacoviello, F. Temperature Influence on Brake Pad Friction Coefficient modelization. *Materials* **2024**, *17* (1), 189. DOI:10.3390/ma17010189.
3. Brembo. Here's why carbon is unavoidable for the discs and brake pads in the top competitions. https://www.brembo.com/en/news-archive/carbon-discs-and-brake-pads-in-the-top-competitions (accessed 2026-02-12).
4. Brembo North America. BREMBO NAMED IZOD INDYCAR SERIES' EXCLUSIVE BRAKE SYSTEMS SUPPLIER. https://www.brembo.com/jp/ComunicatiStampa/2011/Brembo%20Brakes%20Announcement%20DEF.pdf (accessed 2026-02-15).
5. Devi, G. R.; Rao, K. R. Carbon-Carbon Composites -An Overview. *defense Science Journal* **1993**, *43* (4), 369–383. DOI:10.14429/dsj.43.4291.
6. Alnaqi, A. A.; Kosarieh, S.; Barton, D. C.; Brooks, P. C.; Shrestha, S. Material characterization of Lightweight Disc Brake Rotors. *Proceedings of the Institution of Mechanical Engineers, Part L: Journal of Materials Design and Applications* **2018**, *232* (7), 555–565. DOI:10.1177/1464420716638683.
7. Hahn, D. W.; Ozisik, M. N. *Heat Conduction*, 3rd ed.; John Wiley & Sons, Inc.: Hoboken, NJ, **2012**.
8. Talati, F.; Jalalifar, S. Analysis of Heat Conduction in a Disk Brake System. Heat and Mass Transfer 2009, 45, 1047–1059. DOI:10.1007/s00231-009-0476-y.
9. Aijaz, M.; Dar, J. G.; Almanjahie, I. M.; Sajjad, I. Mathematical Simulation and Numerical Computation of the Temperature Profiles in the Peripherals of Human Brain during the Tepid Sponge Treatment to Fever. *Computational and Mathematical Methods in Medicine* **2022**, *2022*, 2658095. DOI:10.1155/2022/2658095.
10. Crank, J. *The Mathematics of Diffusion*, 2nd ed.; Clarendon Press: Oxford, **1975**.
11. Carslaw, H. S. *Introduction to the Mathematical Theory of the Conduction of Heat in Solids*, 2nd ed.; Macmillan and Co., Ltd.: London, **1921**.
12. Spulber, C.; Voloacă, S. Experimental Determination of Emissivity Variation of a Brake Disc by Image Analysis. *Annals of the Academy of Romanian Scientists, Series on Science and Technology of Information* **2010**, *3* (2), 61–71. https://aos.ro/wp-content/anale/IFVol3Nr2Art.4.pdf
13. Atkinson, K.; Han, W.; Stewart, D. *Numerical Solution of Ordinary Differential Equations*; John Wiley & Sons, Inc.: Hoboken, NJ, **2008**.
14. Li, Z.; Qiao, Z.; Tang, T. *Numerical Solution of Differential Equations: Introduction to Finite Difference and Finite Element Methods*; Cambridge University Press: Cambridge, **2018**. DOI:10.1017/9781316678725; 34.
15. Langtangen, H. P.; Linge, S. *Finite Difference Computing with PDEs - A Modern Software Approach*; **2016**. https://hplgit.github.io/fdm-book/doc/pub/book/pdf/fdm-book-4print.pdf (accessed 2026-02-19); 150–160.
16. Luo, X.; Zhuang, C.; Fan, S.; Haugen, N. E. L. Study of Rotating Axisymmetric Flow Using the PISO Algorithm. *Int. J. Numer. Methods Fluids* **2016**, *82*, 708–717. DOI:10.1002/fld.4237; 708–717
17. Strikwerda, J. C. *Finite Difference Schemes and Partial Differential Equations*, 2nd ed.; Society for Industrial and Applied Mathematics (SIAM): Philadelphia, PA, 2004
18. Haque, M.; Akter, S.; Mojumder, S. Numerical Simulation of Convective Heat Transfer for Three-Dimensional Lid-Driven Cavity Flow. *J. Appl. Math. Phys.* **2025**, *13*, 2234–2244. DOI:10.4236/jamp.2025.137127; 2234–2244
19. Mobedi, M.; Barisik, M.; Nakayama, A. Heat and Fluid Flow of Gases in Porous Media with Micropores: Slip Flow Regime. In *Microscale and Nanoscale Heat Transfer: Analysis, Design, and Applications*; Rebay, M.; Kakac, S.; Cotta, R., Eds.; CRC Press: Boca Raton, FL, 2016. DOI:10.1201/b19261; 409–411
20. Tian, T.; Yang, S.; Høyer, J. L.; Nielsen-Englyst, P.; Singha, S. Cooler Arctic Surface Temperatures Simulated by Climate Models Are Closer to Satellite-Based Data than the ERA5 Reanalysis. *Commun. Earth Environ*. **2024**, *5*, 111. DOI:10.1038/s43247-024-01276-z; 5.
21. Mancias, J. D.; Riccardi, M. R.; Berta, M. E. Steady-State Short Crack Propagation in WC–Co Hardmetals. *Materials* **2024**, *17* (10), 2344. DOI:10.3390/ma17102344; 3.
22. Yuan, G.; Li, X.; Dong, Z.; Xiong, X.; Rand, B.; Cui, Z.; Cong, Y.; Zhang, J.; Li, Y.; Zhang, Z.; Wang, J. Pitch-Based Ribbon-Shaped Carbon-Fiber-Reinforced One-Dimensional Carbon/Carbon Composites with Ultrahigh Thermal Conductivity. *Carbon* **2014**, *68*, 413–425. DOI:10.1016/j.carbon.2013.11.018; 7.
23. Mitsubishi Chemical Corporation. DIALEAD Composite: Carbon/Carbon Composites and CMC (Ceramic Matrix Composite) Data Sheet; Mitsubishi Chemical Corporation, 2022. https://www.m-chemical.co.jp/carbon-fiber/pdf/dialead/CC%20CMC.pdf (accessed 2026-02-19); 1.
24. Fédération Internationale de l'Automobile. *2014 Formula One Technical Regulations*; Fédération Internationale de l'Automobile: 2011. https://argent.fia.com/web/fia-public.nsf/A0425C3A0A7D69C0C12578D3002EBECA/$FILE/2014_F1_TECHNICAL_REGULATIONS_-_Published_on_20.07.pdf (accessed 2026-02-15).
25. Octane Photographic Ltd. *2014 German GP F1 Sunday Pitlane and Buildup* (Photo NG/66811); Octane Photos: 2014. https://octanephotos.co.uk/2014/07/20/2014-german-gp-f1-sunday-pitlane-and-buildup/#!photo/NG/66811 (accessed 2026-02-12).
26. Brembo. *20 Years of F1 Brake Discs: Brembo's Transformation*; Brembo News Archive. https://www.brembo.com/en/news-archive/20-years-f1-brake-discs (accessed 2026-02-12).
27. Roccon, A. Boiling Heat Transfer by Phase-Field Method. *Acta Mech*. **2025**, *236*, 5623–5638. DOI: 10.1007/s00707-024-04122-7.
28. Mobedi, M.; Barisik, M.; Nakayama, A. Heat and Fluid Flow of Gases in Porous Media with Micropores: Slip Flow Regime. In *Microscale and Nanoscale Heat Transfer: Analysis, Design, and Applications*; Taylor & Francis Group: 2016. https://minaeng.iyte.edu.tr/wp-content/uploads/sites/250/2021/06/Chapter-16.-Heat-and-Fluid-Flow-of-Gases-in-Porous-Media-with-Micropores-Slip-Flow-Regime.pdf (accessed 2026-02-12)
29. Langtangen, H. P.; Linge, S. *Finite Difference Computing with PDEs: A Modern Software Approach*; Springer: 2017. DOI: 10.1007/978-3-319-55456-3.

DOI: 10.36838/IJHSR820.18

30. Papac, J.; Gibou, F.; Ratsch, C. Efficient Symmetric Discretization for the Poisson, Heat and Stefan-Type Problems with Robin Boundary Conditions. *J. Comput. Phys.* **2010**, *229*, 875–889. DOI: 10.1016/j.jcp.2009.10.017.
31. Luo, K.; Zhuang, Z.; Fan, J.; Haugen, N. E. L. A Ghost-Cell Immersed Boundary Method for Simulations of Heat Transfer in Compressible Flows under Different Boundary Conditions. *Int. J. Heat Mass Transfer* **2016**, *92*, 708–717. DOI: 10.1016/j.ijheatmasstransfer.2015.09.024.
32. Formula 1. *Italy 2014 - Nico Rosberg onboard lap*. https://www.formula1.com/en/video/italy-2014-nico-rosberg-onboard-lap.6061004597001 (accessed 2026-02-13).
33. Qiu, W.; Samruaisin, P.; Chuwattanakul, V.; Maruyama, N.; Hirota, M.; Eiamsa-ard, S. Evaluation of heat transfer performance of a heat exchanger tube mounted with an I-rib twisted tape and twisted winglets. *Scientific Reports* **2025**, *15*, 37047. DOI:10.1038/s41598-025-20806-z.
34. Causon, D. M.; Mingham, C. G. *Introductory Finite Difference Methods for PDEs*; Ventus Publishing ApS, 2010.
35. Kreyszig, E. Advanced Engineering Mathematics, 9th ed.; John Wiley & Sons: Hoboken, NJ, 2006; 593–594
36. Tsega, E. G.; Rao, D. J. Numerical Solution of Three-Dimensional Transient Heat Conduction Equation in Cylindrical Coordinates. Journal of Applied Mathematics 2022, 1993151. DOI:10.1155/2022/1993151; 1–2, 5
37. Blundell, S. J.; Blundell, K. M. Concepts in Thermal Physics, 2nd ed.; Oxford University Press: Oxford, U.K., 2010; 91–92
38. Dellar, P. J. An interpretation and derivation of the lattice Boltzmann method using Strang splitting. *Computers & Mathematics with Applications* **2013**, *65* (2), 129–141. DOI:10.1016/j.camwa.2011.08.047.
39. Ohlhorst, C. W.; Vaughn, W. L.; Tsou, P. Carbon-Carbon Advanced Material for Aircraft Brakes. NASA Technical Memorandum 4787; NASA Langley Research Center: Hampton, VA, 1997; 2–3. https://ntrs.nasa.gov/api/citations/19970041399/downloads/19970041399.pdf (accessed 2026-02-19).

Authors

Max Hong is a student at ACG Parnell College who enjoys physics and mathematics, and has a strong passion for motorsport. He intends to major in mechanical engineering at university and is particularly interested in applying mathematical modeling and simulation to real-world engineering problems.

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

Special Functions: Gamma and Beta Functions - Theory and Applications

Momo Kikuchi

College Alpin International Beau Soleil, Route du Village 1, 1884 Villars-sur-Ollon, Switzerland; momo.kikuchi.12@gmail.com

ABSTRACT: This article investigates the Gamma and Beta functions, two key special functions in mathematics and mathematical physics. It begins by presenting their integral definitions, basic properties, and their application, including the connection between the Gamma function and factorials and the relationship between the Gamma and Beta functions. The applications are also considered, especially in physics and applied mathematics, where these functions appear in problems involving the ideal classical gas, the harmonic oscillator, probability distributions, and integration formulas. This article highlights the importance of the Gamma and Beta functions as useful tools for solving mathematical problems in science and engineering.

KEYWORDS: Mathematical Physics, Gamma Function, Beta Function, Ideal Classical Gas, Harmonic Oscillator, Integration.

Introduction

Mathematics is a fundamental study that supports many fields, including physics, chemistry, natural science, engineering, economics, and architecture. For example, it provides tools for analyzing patterns, modeling real-world phenomena by analyzing statistically, and predicting future events. Although mathematics may seem separate from everyday life, the base of modern society is based on mathematical ideas.

A function is one of the key concepts in mathematics, and some advanced functions appear frequently in mathematics, physics, and engineering, called special functions. For example, the gamma function, beta function, Bessel functions, Legendre polynomials, and spherical functions are special functions. Special functions help solve differential equations, evaluate difficult integrals, and model complex systems from their properties. For instance, the recurrence relation of the Gamma function allows factorials with a continuous number, which can be used to model the time until a specific number of events occur.

This article focuses on the gamma function and the beta function. The gamma function, written as Γ(z), extends the idea of factorials beyond positive integers. It is often used in calculus, probability, statistics, physics, and engineering. For example, it is used in the gamma distribution, chi-squared distribution, improper integrals, differential equations, and formulas involving higher-dimensional geometry.[1]

The beta function is another important special function defined by an integral. It is closely related to the gamma function and has the symmetry property B(x, y) = B(y, x). It is used in probability theory, especially in defining the beta distribution, and it also appears in physics and engineering problems involving trigonometric and weighted integrals.

Together, the gamma and beta functions allow many complicated expressions and integrals to be written in simpler closed forms. They are widely used in theoretical mathematics, physics, statistics, and scientific computation. This article will introduce their definitions, basic properties, important identities, and applications, including Euler's reflection formula and examples from quantum physics, statistical mechanics, and integral evaluation.

Results and Discussion

Properties of the Gamma Function:

Definition:

The Gamma function is denoted as Γ(z), and this is a mathematical function that generalizes the factorial function to complex and real number arguments (except for non-positive integers, where it is undefined). It is defined as:

$$\Gamma(z) = \int_0^\infty t^{z-1}e^{-t}\,dt$$

for complex numbers z with a positive real part.

Recurrence relation:

The recurrence relation shows the value of the function at a given number using the value at the previous number. This relation shows that the value of the gamma function at z + 1 is equal to z multiplied by the value of the gamma function at z, which is related to how factorials are related to each other.[3,4]

When z>0

$$\Gamma(z) = (z-1)\Gamma(z-1) \quad (1)$$

The Gamma function at z equals z – 1 times the Gamma function at z – 1.

Proof:

Solve $\Gamma(z) = \int_0^\infty x^{z-1}e^{-x}\,dx$

$$\Gamma(z) = \int_0^\infty x^{z-1}e^{-x}\,dx$$

$$= [x^{z-1}(-e^{-x})]_0^\infty + \int_0^\infty (z-1)x^{z-2}e^{-x}\,dx$$

Evaluating $[x^{z-1}(-e^{-x})]_0^\infty = 0$

$$x^{z-1}(-e^{-x}) = -\frac{x^{z-1}}{e^x}$$

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.19

In this expression, we assume that z > 1; therefore, the entire expression approaches zero.

Now we are left with the following:

$$\Gamma(z+1) = [-x^z e^{-x}]_0^\infty + z\int_0^\infty x^{z-1}e^{-x}\,dx$$

$$= z\int_0^\infty x^{z-1}e^{-x}\,dx$$

$$= z\Gamma(z)$$

From these proofs, we conclude that:

$$\Gamma(z)=(z-1)\Gamma(z-1)$$

Relation to factorial:

The Gamma function generalizes the factorial function to all complex numbers except for negative integers. For every positive whole number n, the value of the gamma function at n is equal to the factorial value of n – 1.

$$\Gamma(n)=\Gamma(n-1)!$$

The Gamma function at n is equal to the factorial n – 1.[3,4]

Proof:

The recurrence relation of the Gamma function shows why the Gamma function is connected to the factorial function. Since equation (1),

$$\Gamma(z)=(z-1)\Gamma(z-1)$$

has already been proved, we can use it to show that the Gamma function gives the same values as factorials when the input is a positive integer. First, we check the base case:

$$\Gamma(1) = \int_0^\infty e^{-t}\,dt = 1$$

Also,

$$0!=1$$

Therefore,

$$\Gamma(1)=0!$$

Now, using the recurrence relation, we obtain

$$\Gamma(z)=(z-1)(z-2)(z-3)\ldots\Gamma(1)$$

Since $\Gamma(1)=1$, we obtain

$$\Gamma(z)=(z-1!)$$

From these proofs, we conclude that:

$$\Gamma(n)=\Gamma(n-1)!$$

This means that the Gamma function gives the same values as factorials for positive integer inputs, while also allowing similar calculations for non-integer and complex values.[3]

This can be used in probability and statistics. For example, the Gamma function is used in the gamma distribution.[3] This distribution is commonly used to model waiting times like the time until a machine fails or the time between random events.[5] In these situations, the shape of the distribution does not always have to be a whole number. It may be a decimal value, such as 2.5. By using this factorial relation of the gamma function, it is possible to calculate the decimal value of a factorial.

Special value:

Some special values of the Gamma function are important, allowing us to connect with π. The Gamma function at 1/2 is equal to $\sqrt{\pi}$.

$$\Gamma(1/2) = \pi$$

In this proof, we use the Gaussian integral, I , which is defined as

$$I = \int_{-\infty}^{\infty} e^{-x^2}\,dx$$

Since this integral is difficult to evaluate directly, we square it:

$$I^2 = \left(\int_{-\infty}^{\infty} e^{-x^2}\,dx\right)\left(\int_{-\infty}^{\infty} e^{-y^2}\,dy\right)$$

This variable y is only used to create a copy of I. Since they are both integrals of the same product, their product is still I^2.

I^2 can also be written as:

$$I^2 = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-(x^2+y^2)}\,dxdy$$

To simplify, we use polar coordinates:

$$x^2 + y^2 = r^2$$

Therefore,

$$I^2 = \int_0^{2\pi}\int_0^{\infty} e^{-r^2}\,r\,drd\theta$$

This integral can be separated:

$$I^2 = \left(\int_0^{2\pi} d\theta\right)\left(\int_0^{\infty} re^{-r^2}\,dr\right)$$

These two integrals can be calculated as:

$\int_0^{2\pi} d\theta = 2\pi,$

and

$$\int_0^{\infty} re^{-r^2}dr = \frac{1}{2}$$

Therefore,

$$I^2 = 2\pi \times \frac{1}{2} = \pi$$

Since I>0,

$$I = \sqrt{\pi}$$

From these proofs, we conclude that:

$$\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}$$

This special value of the Gamma function gives exact values, connecting the Gamma function to. This value is also used to calculate other half-integer values.

For example,

$$\Gamma\left(\frac{3}{2}\right) = \frac{1}{2}\Gamma\left(\frac{1}{2}\right) = \frac{\sqrt{\pi}}{2}$$

They allow us to simplify difficult expressions instead of leaving an answer in complicated form.

Euler's Reflection Formula:

Euler's reflection formula is an identity that relates the Gamma function at z and 1-z.[3,4] It is especially important in complex analysis and the theory of special functions.[4]

For complex numbers z where z is not an integer:

$$\Gamma(z)\Gamma(1-z) = \frac{\pi}{sin\ sin\ (\pi z)}$$

Proof:

We begin with Euler's infinite product representation:

$$\frac{1}{\Gamma(z)} = ze^{yz}\prod_{n=1}^{\infty}\left(\left(1+\frac{z}{n}\right)e^{-\frac{z}{n}}\right)$$

where γ is the Euler-Mascheroni constant.

Now replace z by $-z$. We obtain

$$\frac{1}{\Gamma(-z)} = (-z)e^{-\gamma z}\prod_{n=1}^{\infty}\left(\left(1-\frac{z}{n}\right)e^{\frac{z}{n}}\right)$$

Next, we multiply these two expressions:

$$\frac{1}{\Gamma(z)\Gamma(-z)} = \left[ze^{z\gamma\prod_{n=1}^{\infty}\left(1+\frac{z}{n}\right)e^{-\frac{z}{n}}}\right]\left[(-z)e^{-\gamma z\prod_{n=1}^{\infty}\left(1-\frac{z}{n}\right)e^{\frac{z}{n}}}\right]$$

The exponential factors cancel:

$$e^{\gamma z}e^{-\gamma z} = 1$$

$$e^{-\frac{z}{n}}e^{\frac{z}{n}} = 1$$

Therefore,

$$\frac{1}{\Gamma(z)\Gamma(-z)} = -z^{2\prod_{n=1}^{\infty}\left(1+\frac{z}{n}\right)\left(1-\frac{z}{n}\right)}$$

Since $(1+a)(1-a)=1-a^2$,

$$\frac{1}{\Gamma(z)\Gamma(-z)} = -z^2\prod_{n=1}^{\infty}\left(\left(1-\frac{z^2}{n^2}\right)\right) \quad (2)$$

Now we use the infinite product expansion of the sine function:

$$sin\,sin\,(\pi z) = \pi z\prod_{n=1}^{\infty}\left(\left(1-\frac{z^2}{n^2}\right)\right)$$

Therefore,

$$\prod_{n=1}^{\infty}\left(\left(1-\frac{z^2}{n^2}\right)\right) = \frac{sin\,sin\,(\pi z)}{\pi z}$$

We substitute this into (2),

$$\frac{1}{\Gamma(z)\Gamma(-z)} = -z^2\times\frac{sin\,sin\,(\pi z)}{\pi z}$$

We cancel out z,

$$\frac{1}{\Gamma(z)\Gamma(-z)} = -\frac{zsinsin\,(\pi z)}{\pi} \quad (3)$$

Now we use the recurrence relation (1),

$$\Gamma(z)=(z-1)\Gamma(z-1)$$

Replace z by $1-z$,

$$\Gamma(1-z)=(-z)\Gamma(-zz)$$

Multiply both sides by $\Gamma(z)$,

$$\Gamma(z)\Gamma(1-z)=-z\Gamma(z)\Gamma(-z)$$

Substitute for (3),

$$\Gamma(z)\Gamma(-z) = -\frac{\pi}{z\,sin\,sin\,(\pi z)}$$

Then,

$$\Gamma(z)\Gamma(1-z) = -z\left(-\frac{\pi}{z\,sin\,sin\,(\pi z)}\right)$$

Therefore,

$$\Gamma(z)\Gamma(1-z) = \frac{\pi}{sin\,sin\,(\pi z)}$$

From these proofs, we conclude that

$$\Gamma(z)\Gamma(1-z) = \frac{\pi}{sin\,sin\,(\pi z)}$$

Euler's reflection formula gives a relationship between $\Gamma(z)$ and $\Gamma(1-z)$. This allows difficult Gamma function values to be expressed in terms of trigonometric functions. Special values such as $\Gamma\left(\frac{1}{2}\right)=\sqrt{\pi}$ and for the Gamma function at negative non-integer values can be calculated using Euler's reflection formula.

Applications of the Gamma Function:

Quantum Mechanics: Quantum Harmonic Oscillator:

The harmonic Oscillator is one of the important models of physics. Any mass experiencing a restoring force near a stable equilibrium behaves like a harmonic oscillator for small vibrations, and because of this, harmonic oscillators appear throughout nature and form the basis of many human-made technologies, including clocks and radio circuits. In quantum physics, energy comes in fixed, discrete levels rather than a smooth range in the system of a harmonic oscillator.[6] This simple model turns out to be very powerful because it helps describe vibrations of atoms, molecules, and even the behavior of light in a cavity. We will study in this section how the Gamma function appears in the quantum harmonic oscillator.

A Gaussian moment is the expected value of the power of a variable in a Gaussian distribution. The Gaussian distribution (also called the normal distribution) is a type of symmetric bell-shaped probability distribution. In one-dimensional quantum harmonic oscillators, the wave function of the ground state is Gaussian, so these moments naturally appear.[6] Therefore, quantities such as $\langle x^{2k}\rangle$ can be directly reduced to Gaussian moments.[7]

Consider the general Gaussian moment:

$$I_n = \int_0^{\infty} x^n e^{-x^2}\,dx,\ (n>1)$$

Let $u=x^2$, $du=2xdx$,

$$x^2\,dx = \frac{1}{2}u^{\frac{n-1}{2}-1}e^{-u}\,du$$

Substitute:

$$I_n = \frac{1}{2}\int_0^{\infty} u^{\frac{n-1}{2}-1}e^{-u}\,du$$

From the definition of the Gamma function,

$$I_n = \frac{1}{2}\Gamma\left(\frac{n+1}{2}\right)$$

Therefore,

$$\int_0^{\infty} x^n e^{-x^2}\,dx = \frac{1}{2}\Gamma\left(\frac{n+1}{2}\right)$$

By symmetry,

$$\int_{-\infty}^{\infty} x^n e^{-x^2}\,dx = \begin{cases}0, n\ odd\\ \Gamma\left(\frac{n+1}{2}\right), n\ even\end{cases}$$

Example 1: Normalization of the ground-state wave function:

In quantum mechanics, a wave function must be normalized, meaning that the total probability of finding the particle must be equal to 1. Therefore, the ground-state wave function must satisfy

$$\int_{-\infty}^{\infty} |\Psi_0(x)|^2\,dx = 1$$

We substitute

$$\Psi_0(x) = Ae^{-\frac{\alpha x^2}{2}}$$

obtaining

$$|\Psi_0(x)|^2 = A^2e^{-\alpha x^2}$$

DOI: 10.36838/IJHSR820.19

Therefore, the normalization condition becomes

$$A^2 \int_{-\infty}^{\infty} e^{-\alpha x^2}\, dx = 1$$

The integral $\int_{-\infty}^{\infty} e^{-\alpha x^2}\, dx$ is a Gaussian integral. Using the Gamma function, it can be simplified as:

$$\int_{-\infty}^{\infty} e^{-\alpha x^2}\, dx = \sqrt{\frac{\pi}{\alpha}}$$

Therefore,

$$A^2 \sqrt{\frac{\pi}{\alpha}} = 1$$

Solving for *A*, we get

$$A = \left(\frac{\alpha}{\pi}\right)^{\frac{1}{4}}$$

Therefore, the normalized ground state wave function is

$$\Psi_0(x) = \left(\frac{\alpha}{\pi}\right)^{\frac{1}{4}} e^{-\frac{\alpha x^2}{2}}$$

This shows that the Gamma function is useful when determining the correct normalization constant for the quantum harmonic oscillator.

Example 2: Expectation value of x in the ground state:

The expectation value of position is defined by

$$\langle x \rangle = \int_{-\infty}^{\infty} x|\Psi_0(x)|^2\, dx$$

Substituting the normalized probability density gives

$$\langle x \rangle = A^2 \int_{-\infty}^{\infty} x e^{-\alpha x^2}\, dx$$

The function $xe^{-\alpha x2}$ is an odd function since replacing x by –x changes its sign. Since the integral is taken over the interval (-∞, ∞), the positive and negative contributions cancel each other.

Therefore,

$$\langle x \rangle = 0$$

This shows that the average position of the particle is zero in the ground state, meaning that the particle has no preferred direction from the center of the oscillator.

Ideal Classical Gas and the Equipartition Theorem

An ideal gas is a model for a gas with a few assumptions, like negligible particle volume, no intermolecular forces, perfectly elastic collisions, and random motion.[8] This model follows the ideal gas law, PV=nRT. It is easy to handle in statistical mechanics and useful when calculating pressure, volume, mole, and temperature. In reality, gas particles do interact, have a measurable volume, and can liquify or solidify at low temperatures and high pressure.

In classical statistical mechanics, the equipartition theorem gives a direct link between a system's temperature and its average energy.[9] It is often referred to as the law of equipartition. The basic idea is that, at thermal equilibrium, energy is shared evenly among all available forms, so that each degree of freedom carries the same average kinetic energy.

For a gas particle moving in three dimensions, the particle can move in 3 direction of: x, y, and z. Therefore, it has three translational degrees of freedom. Each direction contributes:

$$\frac{1}{2} k_B T$$

So the average kinetic energy of one particle is:

$$3 \times \frac{1}{2} k_B T = \frac{3}{2} k_B T$$

If a gas contains *N* particles, the kinetic energy will be

$$U = \frac{3}{2} N k_B T$$

This formula shows that the internal energy of a classical ideal gas is proportional to the number of particles and the temperature. As the temperature increases, the internal energy increases.

This result can be calculated using the Gamma function in a more advanced derivation. The average energy of gas particles can be calculated using integrals over momentum in statistical mechanics. These integrals have the form:

$$\int_0^{\infty} x^{n-1} e^{-\alpha x^2}\, dx$$

This can be written using the Gamma function:

$$\int_0^{\infty} x^{n-1} e^{-\alpha x^2}\, dx = \frac{1}{2} \alpha^{-\frac{n}{2}} \Gamma\left(\frac{n}{2}\right)$$

The gamma function gives a more advanced way to obtain the same result through statistical mechanics, giving the internal energy formula,

$$U = \frac{3}{2} N k_B T$$

Properties of Beta Function:

Definition:

The Beta function, denoted as B(x, y), is a special function defined for real numbers

x, y > 0. It is given by the integral:

$$B(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1}\, dt$$

It is also known as the Euler integral of the first kind.

Symmetry Property:

The symmetry property states that the Beta function is unchanged if we swap its two input values:[10]

$$B(m,n) = B(n,m)$$

This means the Beta function treats its two arguments equally. Geometrically, this reflects the symmetry of the integrand when the variable t is replaced with 1 t. For example, B(2, 3) is equal to B(3, 2).

Proof:

We start from the definition

$$B(m, n) = \int_0^1 t^{m-1}(1-t)^{n-1}\, dt$$

Now make the substitution t = 1 – u. Then

$$dt = -du$$

and when t goes from 0 to 1, u goes from 1 to 0.

So the integral becomes

$$B(m,n) = \int_{t=0}^{1} t^{m-1}(1-t)^{n-1}\,dt = \int_{u=1}^{0} (1-u)^{m-1}u^{n-1}(-du)$$

Reversing the limits of the integral gives

$$B(m,n) = \int_{0}^{1} u^{n-1}(1-u)^{m-1}\,du$$

which is the definition of $B(n,m)$, therefore,

$$B(m,n)=B(n,m)$$

From this proof, we conclude that swapping the two inputs of the Beta function does not change its value. This symmetry often simplifies calculations and formulas. For example, in statistics, the Beta distribution has two "shape" parameters, and its formula does not change if those two parameters are swapped. In integrals, this symmetry means you can swap the roles of the two powers and still get the same result.

Relation with Gamma Function:

This property connects the Beta function to the Gamma function, Γ(z), which generalizes the factorial. The relationship is:

$$B(x,y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)} \quad (4)$$

Recall that for natural numbers:

$$\Gamma(n) = (n-1)!$$

So, the Beta function can also be written using factorials when x and y are whole numbers. This formula is extremely useful in simplifying expressions involving Beta and Gamma functions.[3]

Proof:

To start with, we will rewrite (4) as:

$$\Gamma(x)\Gamma(y)=\Gamma(x+y)B(x,y)$$

Next, we will look at the Gamma function definitions,

$$\Gamma(x) = \int_{0}^{\infty} u^{x-1}e^{-u}\,du$$

and

$$\Gamma(y) = \int_{0}^{\infty} v^{y-1}e^{-v}\,dv$$

Multiplying these two integrals will give

$$\Gamma(x)\Gamma(y) = \int_{0}^{\infty}\int_{0}^{\infty} u^{x-1}v^{y-1}e^{-(u+v)}\,dudv \quad (5)$$

Simplify by

$$r = u + v$$

and

$$z = \frac{u}{u+v}$$

The Jacobian for this substitution is

$$du\,dv = r\,dr\,dz$$

After substitution, the integral (5) becomes

$$\int_{0}^{1}\int_{0}^{\infty} (rz)^{x-1}[r(1-z)]^{y-1}e^{-r}\,r\,drdz$$

Now separate the powers of r, z, and 1-z

$$(rz)^{x-1}[r(1-z)]^{y-1}r = r^{x+y+1}z^{x-1}(1-z)^{y-1}$$

Therefore, the integral becomes

$$\int_{0}^{1}\int_{0}^{\infty} r^{x+y+1}e^{-r}z^{x-1}(1-z)^{y-1}\,drdz$$

This can be written as

$$\Gamma(x)\Gamma(y) = \left(\int_{0}^{\infty} r^{x+y-1}e^{-r}\,dr\right)\left(\int_{0}^{1} z^{x-1}(1-z)^{y-1}\,dz\right)$$

The first part of this integral is the Gamma function:

$$\int_{0}^{\infty} r^{x+y-1}e^{-r}\,dr = \Gamma(x+y)$$

The second part of this integral is the Beta Function:

$$\int_{0}^{\infty} z^{x-1}(1-z)^{y-1}\,dz = B(x,y)$$

Therefore,

$$B(x,y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$$

which equals

$$B(x,y+1) = \frac{y}{x+y}B(x,y)$$

This relates the beta function to the gamma function. Applying this allows us to calculate beta expressions using given gamma values. For example, it ensures that the sum of all possible outcomes in a beta probability equals 1. In calculus, it helps calculation by simplifying beta integrals into gamma values.

Reduction Formula:

The reduction formula is a useful identity that allows us to express a Beta function with one of its parameters increased by 1 in terms of the original Beta function. The Beta function can be difficult to compute directly for all values. But with this formula, we can reduce a harder case to an easier one. We want to show:

$$B(x,y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$$

Proof:

Using the relation between the Beta and Gamma functions,

$$B(x,y+1) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y+1)}$$

substitute y+1

$$B(x,y+1) = \frac{\Gamma(x)y\Gamma(y)}{(x+y)\Gamma(x+y)}$$

Using the recurrence relation of the Gamma function,

$$\Gamma(y+1)=y\Gamma(y)$$

and

$$\Gamma(x+y+1) = (x+y)\Gamma(x+y)$$

Therefore,

$$B(x,y+1) = \frac{\Gamma(x)y\Gamma(y)}{(x+y)\Gamma(x+y)}$$

DOI: 10.36838/IJHSR820.19

This can be written as

$$B(x, y+1) = \frac{y}{x+y} \times \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$$

Going back to the relation between the Beta and Gamma functions,

$$\frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)} = B(x, y)$$

Therefore,

$$B(x, y+1) = \frac{y}{x+y} B(x, y)$$

These formulas let us calculate B(x,y) by reducing x or y by 1. This is useful when evaluating integrals of polynomials or trigonometric powers. Each application of the reduction formula gets rid of one power. For example, it is used in integration techniques when reducing $\int \theta \, d\theta$ or $\int \theta \, d\theta$ one power at a time.

Trigonometric Representation:

This formula gives a new way to compute the Beta function, not with powers of t, but with powers of sin θ and cos θ, where θ ranges from 0 to π/2.[10] We show that

$$B(x, y) = 2 \int_0^{\frac{\pi}{2}} (sin\theta)^{2x-1} (cos\theta)^{2y-1} \, d\theta$$

Proof:

We begin with the definition of the Beta function:

$$B(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} \, dt \quad (5)$$

We are going to change the variable of integration from t to θ . Let,

$$t = \theta$$

First, we find how the limit changes from this substitution:

$$t = 0 \Rightarrow \theta = 0 \Rightarrow \theta = 0$$
$$t = 1 \Rightarrow \theta = 1 \Rightarrow \theta = \pi/2$$

Now, differentiate:

$$t = \theta$$
$$dt = 2sin\theta \, cos\theta \, d\theta$$

Now we substitute into the original integral (5):

$$B(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} \, dt$$

Since

$$t^{x+1} = (\theta)^{x-1} = (sin\theta)^{2x-2}$$

and similarly,

$$(1-t)^{y-1} = (\theta)^{y-1} = (cos\theta)^{2y-2}$$

Therefore

$$B(x, y) = \int_0^{\frac{\pi}{2}} (sin\theta)^{2x-2} (cos\theta)^{2y-2} \, 2sin\theta cos\theta \, d\theta$$

Now, we can simplify by combining powers:

$$(sin\theta)^{2x-2} \times sin\theta = (sin\theta)^{2x-1}$$
$$(cos\theta)^{2x-2} \times cos\theta = (cos\theta)^{2x-1}$$

Therefore,

$$B(x, y) = 2 \int_0^{\frac{\pi}{2}} (sin\theta)^{2x-1} (cos\theta)^{2y-1} \, d\theta$$

This is the trigonometric representation of the Beta function. It is especially useful for integrals that involve powers of sine and cosine. Such integrals appear when finding areas under curves or volumes of solids with circular symmetry. In physics and engineering, this representation also appears when studying sinusoidal signals, such as sound waves, alternating electric currents, and vibrations. The Beta function helps turn these complicated trigonometric integrals into simpler expressions.

■ Conclusion

This article has summarised the key properties of the Gamma function and Beta function. It covers the definition of the Gamma function to start with, then proves the recurrence relation, which shows that the value of the Gamma function at z + 1 is equal to z multiplied by the value of the Gamma function at z.[3] From this relation, we covered the factorial relation. We also discussed the special values, including the calculation of $\Gamma(1/2) = \sqrt{\pi}$. In addition, we proved Euler's Reflection Formula, which is an identity that relates Γ(z) and Γ(1 − z) through the sine function.

For the Beta function, we showed its symmetry property. This means that the order of the arguments does not matter. We then proved its close relationship with the Gamma function. This connection highlights how the Beta function can be expressed using Gamma functions, which can be applied to simplify calculations and solve more complex problems. We also explored the reduction formula, which allows more complicated Beta function expressions to be simplified. The trigonometric representation, which ties the Beta function to integrals involving sine and cosine, was also proved. Finally, we covered special cases such as B(1, n), which simplify simple fractions.

We also introduced Gaussian moments and the one-dimensional quantum harmonic oscillator to show how these functions are applied in physics, particularly for describing distributions and quantum systems. For example, the Gamma function helps calculate probabilities and energies in these contexts. In addition, we discussed the classical ideal gas in three dimensions and its internal energy. We have explained how these mathematical tools can be applied to real physical systems in thermodynamics. Also, the ideal classical gas, where we looked at the internal energy of a classical ideal gas in 3 dimensions, can be applied to chemistry, physics, and other various fields.

In conclusion, these properties allow us to calculate more complicated problems more simply. Because of their extremely useful properties, they are applied in various fields such as physics, probability and statistics, calculus, engineering, and many others. The Gamma function and Beta function described in this manuscript are widely used in the real world.

■ Acknowledgments

I would like to thank my mentor, Mir Mehedi Faruk, for his guidance and helpful feedback throughout this research. His support helped me develop a clearer understanding of the

Gamma and Beta functions and their applications in mathematical physics.

References

1. Abramowitz, M. and Stegun, I.A., 1972. Handbook of mathematical functions with formulas, graphs, and mathematical tables. New York: Dover Publications
2. Forster, O., 2015. 9. The Gamma Function. In: Analytic Number Theory. Munich: LMU. [pdf] Available at: https://www.mathematik.uni-muenchen.de/~forster/ v/ann/annth_09.pdf
3. Sebah, P. and Gourdon, X. (2002). *Introduction to the Gamma Function*. Available at: numbers.computation.free.fr
4. Arfken, G.B., Weber, H.J., and Harris, F.E., 2013. Mathematical methods for physicists. 7th ed. Amsterdam: Elsevier
5. Tong, D., 2016. Statistical Physics. University of Cambridge. [online] Available at: https://www.damtp.cam.ac.uk/user/tong/statphys/statphys.pdf
6. Griffiths, D.J. and Schroeter, D.F. (2018). *Introduction to Quantum Mechanics*. 3rd edn. Cambridge: Cambridge University Press
7. Bertlmann, R.A. (2010). Chapter 5: Harmonic Oscillator and Coherent States. University of Vienna. Available at: https://homepage.univie.ac.at/reinhold.bertlmann/pdfs/T2_Skript_Ch_5.pdf
8. Schroeder, D.V. (2021). An Introduction to Thermal Physics. 2nd edn. Oxford: Oxford University Press
9. Lima, J.A.S. and Deppman, A. (2000). 'On the classical energy equipartition theorem', Brazilian Journal of Physics, 30(1), pp. 176–180

Author

Momo Kikuchi is an IB student at College Alpin International Beau Soleil. She is currently studying the application of special functions in physics and special functions involving complex analysis.

DOI: 10.36838/IJHSR820.19

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Developing Next-Generation dCas9-VP64 Systems for Temporal Activation Functionalities

Rohaan Sanyal

Sacred Heart Preparatory, 150 Valparaiso Avenue, Atherton, CA, 94025, USA; rohaansanyal@gmail.com

ABSTRACT: Recent developments have aimed to improve the efficiency and specificity of CRISPRa (activation) and CRISPRi (interference), different complexes involving the dCas9 enzyme, aiming to upregulate or downregulate genes through promoter binding. Allosteric regulation in the arCas9 system enables inducible and temporal control of CRISPR, using small molecules and corresponding ligand receptors. These engineered systems integrate ligand-binding domains (e.g., the estrogen receptor ESR1) directly onto the Cas9 or onto scaffold/effector domains. This permits activation of CRISPR based on the presence of specific signals (4-Hydroxytamoxifen in the case of ESR1), allowing precise ON/OFF switching. This allostery represents an improvement from the original CRISPR-Cas9 systems, which are constitutive in the presence of guide RNA (sgRNA) and often inadvertently create off-target effects. Building on the framework of ligand-based control and dCas9 gene regulation, we theorized a construct containing a dead allosterically regulated Cas9 (darC9) with effector domains. A darC9-effector domain platform could activate genes or inhibit regulators in the presence of ligands, allowing reversible and controlled triggers for various genetic engineering applications. Utilizing a Gibson assembly, we tested the combination of a recipient dCas9-VP64 plasmid with a sourced ESR1 to generate a darC9-VP64 molecule with precise ON/OFF capabilities. We completed an in silico design of a darC9-VP64 tool with theoretical imaging capabilities through GFP. We found that amplification of the 5' Gibson assembly fragment was successful, but we had challenges amplifying the ESR1 and 3' Gibson assembly fragment.

KEYWORDS: Biology, Molecular Biology, arCas9, CRISPR, dCas9.

Introduction

Since its inception, the Clustered Regularly Interspaced Short Palindromic Repeats (CRISPR) gene editing system has revolutionized the field of gene therapy and holds immense potential for many other fields in therapeutics and medicine. The system, popularized in 2012, is capable of locating and making cuts within the nucleus, using the Cas (CRISPR-associated protein) molecule (sourced from Streptococcus pyogenes or related bacteria) and short sequences called single-guide RNAs (sgRNAs). The Cas molecule employs its two nucleases, RuvC and HNH, to cleave the complementary and target strands of DNA, respectively.[1]

Mutations of the Cas molecule at designated sites (D10A and H840A)[2] produce a dead Cas9 (dCas9), a molecule that is able to locate desired loci but is unable to make cuts, minimizing the off-target effects (unintended edits in other loci) of CRISPR. A derivative of this novel process is the CRISPRi and CRISPRa structures, which utilize CRISPR's locating abilities with a dCas9 in conjunction with an activation or deactivation domain to reversibly upregulate or downregulate a gene.

CRISPR interference (CRISPRi) works to downregulate a gene, utilizing a dCas9-sgRNA complex that binds to a complementary strand of a desired region, blocking transcription elongation using effector domains (such as KRAB). Conversely, CRISPR activation (CRISPRa) enables the targeting of any gene using specific and programmable sgRNA sequences to upregulate the selected gene. It localizes itself near a locus of interest and recruits transcription factors using activation effector domains (such as VP64 or p65) that are fused to one or both N- and C- termini of the dCas9. The effector domain recruits endogenous transcriptional activators, making downstream transcription of a protein-coding sequence more efficient, which successfully upregulates a gene.[3]

Advances to the first-generation dCas9-VP64 system form the second generation of CRISPR activation/Cas. The CRISPR SAM system (Synergistic Activation Mediator) combines dCas9 with specifically engineered sgRNAs to substantially boost transcription and activation. The dCas9 is fused with aptamers that bind to MS2 proteins, recruiting additional activation domains. SunTag, another second-generation CRISPRa approach, uses more VP64 per dCas9 with a repeated peptide scaffold, outperforming first-generation activation levels.[4,5]

A further advancement to CRISPR is incorporating inducibility and allosteric regulation. Conditionally activated Cas9s are capable of 25-fold higher specificity in target genomic sites of the human cell than wild-type Cas9 (wtCas9).[6] The first iteration of arCas9 (Oakes *et al.*)[7] fuses a ligand-binding domain (LBD) to a Cas9 to create arCas9. The traditional ligand-binding domain used in arCas9 is an estrogen receptor ligand-binding domain (ER-LBD), usually an ESR1, that binds with 4-Hydroxytamoxifen (4-OHT).[7] Other approaches to conditionally activated Cas involve chemically induced plasmids, or intein-mediated splits in the Cas9 molecule (Table 1).

A Polymerase Chain Reaction (PCR) is a nucleic acid amplification technique that aims to create millions to billions of exact copies of a DNA target based on the number of cycles. Each cycle has a denaturation, annealing, and extension step.[8]

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.20

A Gibson assembly is a cloning procedure used to combine multiple features of different plasmids into one. A desired plasmid is produced by amplifying overlapping fragments (usually with PCR) and combining the amplified fragments with a 5' exonuclease, a DNA polymerase, and a DNA ligase. This allows for the desired sequence to bind with a linearized recipient plasmid at a specified target site.[9] Both of these techniques are essential tools for biomolecular research.

Based on these prior efforts, we theorized the creation of a darC9-VP64 system that would be capable of inducible allosteric activation, permitting reversible upregulation of a target gene. A functional darC9-VP64 tool would be able to temporally activate desired genes after ligand administration, which would be particularly effective for *in vivo* therapeutic applications requiring the minimization of off-target effects.

We developed a system to produce a darC9-VP64 plasmid through an *in silico* Gibson cloning assembly using a host and recipient plasmid to produce ESR1 and dCas9-VP64, respectively. The Gibson cloning assembly involved three fragments: the region from AgeI (the upstream restriction site) to the ERBD insertion site, the ERBD, and the region from the ERBD insertion site to BamHI (the downstream restriction site). Experimentally, we saw successful PCR amplification of the AgeI fragment, but were unsuccessful for the remaining fragments, demonstrating that darC9-VP64 is a feasible tool.

Methods

Theoretical design and engineering review:

A literature review was conducted to assemble various approaches to engineering Cas9 systems. Queries were done in PubMed, using the keywords "inducible", "Cas9", "CRISPRi", and "CRISPRa". Three tables were assembled to summarize findings for each system: Table 3 for inducible approaches, Table 4 for CRISPRi, and Table 5 for CRISPRa, found in Appendix A.

Among the reviewed CRISPRa systems, first-generation dCas9-VP64 provides the weakest activation levels, requiring multiple gRNAs and producing off-target effects. Second-generation approaches such as VPR and SunTag improve activation potency through multiple effector domains, but suffer from variable efficiency. All reviewed systems rely on constitutive promoters, which provide no temporal control over gene activation. This gap motivates the development of a darC9-VP64 tool that can combine the activation functionality of CRISPRa with the allosteric control from a ligand-binding domain.

Plasmid templates:

Two plasmids were used as templates and recipients. dCAS9-VP64_GFP was a gift from Feng Zhang (Addgene plasmid #61422; http://n2t.net/addgene:61422; RRID:Addgene_61422). It was used as a recipient for the ESR1 insertion. pBLO1811_arC9_noNLS_human was a gift from David Savage (Addgene plasmid #74494; http://n2t.net/addgene:74494 ; RRID:Addgene_74494), which was used to source the ESR1.

Plasmid streaking and colony validation:

To conduct the plasmid streaking protocol, LB-agar plates containing 100 µg/mL ampicillin were labeled with the plasmid identification and date. A sterile pipette tip was used to collect cells from the plasmid agar stab, which were then inoculated onto the plate by streaking to establish a dense initial region. Serial dilution streaking was performed by transferring cells with fresh tips in a spiral pattern to promote the isolation of single colonies. The inoculated plates were incubated overnight at 37˚C. Colony growth was assessed the following day, and plates were sealed with Parafilm and stored appropriately.

Thermal cycling:

The T100 BioRad Thermal Cycler was used for all thermal cycling applications during the procedure. The thermocycling conditions were as follows: denaturation at 95˚C for 2 minutes; 33 cycles of 95˚C for 30 seconds, primer annealing temperature for 30 seconds (Table 2), 72˚C for 2 minutes and 15 seconds; followed by a final extension at 72˚C for 5 minutes.

In silico:

Addgene was used for ordering and validating sequences. BLAST was used to validate and characterize plasmids sourced from Addgene, utilizing a tiling approach where sequences were broken up into 500bp chunks for more accurate queries. Benchling was used for sequence alignments to validate the dCas9-VP64 plasmid and the ESR1 plasmid (Figure 1). Benchling was also used for the Gibson assembly (link to final sequence can be found in Appendix B), annotation, and primer design.



Figure 1: Sequence alignment of dCas9-VP64 #61422 and arC9 ESR1 template (see *Plasmid templates*). Black lines at the top represent "votes" (homologous pairs between the two sequences). Red bars above the plasmid annotations represent sections that are "skipped" by the alignment due to a lack of homology. Both Cas9 sections are entirely homologous except for the ERBD in the aligned sequence (bottom) and the D10A and H840A mutations in the template sequence (top).

Validation:

Validation was done by creating digestion mixtures that included designated amounts of CutSmart buffer, restriction enzymes BamHI and AgeI, and colony DNA. The mixture was then incubated at 37˚C for 16 hours. The mixture was then loaded into a gel cast that ran at 180V for 4-5 hours. The gel was then imaged using the Invitrogen iBright 1500 and excised.

Throughout the process, gel electrophoresis was employed to validate digestion and PCR products, and primer design was iteratively refined based on experimental outcomes. The final steps involved isolating validated DNA fragments and preparing them for subsequent cloning procedures. This was done by visualizing the band under the iBright imager and

excising it using a scalpel. The gel was then digested with the Qiagen Gel Isolation Kit and its corresponding protocol.

PCR Setup & Conditions:

Primer working stock was set up as follows: 2.5 µL of Taq buffer (1X); 0.5 µL of dNTPs (200 µM); 1 µL of MgCl2 (2 mM); 5 µL of the FWD and REV primers at 1 µM; 100 ng of gDNA; 0.25 µL of Taq polymerase (1.25 U); and RNase-free water until the stock reached 25 µL.

Results and Discussion

Theorized design:

We presumed that integration of a ligand-binding domain onto a catalytically inactive Cas9 fused to an effector domain enables inducible, precise gene expression or regulation with high temporal control. In the proposed design, binding of the small-molecule ligand 4-hydroxytamoxifen to an estrogen receptor ligand-binding domain (ER-LBD) triggers a conformational switch, shifting the complex to an ON state, and allows it to locate and subsequently recruit transcription factors.

In silico assembly of darC9-VP64 reveals a promising engineering strategy:

To engineer the darCas9-VP64 construct, a systematic in silico workflow was implemented, beginning with a review of plasmid design strategies. Three approaches were considered: 1. cloning VP64 onto a dCas9-ESR1 plasmid,[7] 2. inserting ER-LBD into the dCas9-VP64 backbone, and 3. an approach involving site-directed mutagenesis of arCas9-VP64. Due to available resources and construct feasibility, the second approach was selected, utilizing the dCas9-VP64 plasmid as the recipient and using the arCas9 plasmid for the ER-LBD template. These vectors were chosen due to their compatibility for Gibson cloning, uniform delivery method, and corresponding restriction sites.

Key restriction sites, AgeI and BamHI, were identified as optimal insertion points for ER-LBD based on sequence alignment, and corresponding primers were designed to amplify these regions. The ER-LBD insertion site was identified at 231.[7]

A construct was designed to assemble darC9-VP64 in Benchling. dCas9-VP64 served as a backbone. The host plasmid was then digested at the AgeI and BamHI sites. One primer pair covered the region from AgeI to the ESR1 insertion site at 231, and the other primer pair covered the downstream region from 231 to BamHI. The individual fragments of the assembly are shown in Table 1. The primer pairs for PCR are shown in Table 2.

Based on the construct, a theoretical procedure was created: to assemble darC9-VP64, restriction enzymes would digest the recipient plasmid (dCas9-VP64) at the selected restriction sites (AgeI and BamHI) in order to create an opening for ERBD insertion. After restriction enzyme digestion, each fragment would be amplified with PCR, and then Gibson cloning would attach the fragments to assemble the completed plasmid.

Table 1: List of fragments from *in silico* Gibson assembly on Benchling. Color corresponds to Figure 3. Fragments 2-4 were amplified. "Start" and "End" refer to the base pair coordinates from each fragment's source plasmid.

#	Sequence	Source	Bin	Start	End
1	dCas9-VP64_GFP #61422	#61422	Backbone	8094	2760
2	AgeI to insertion site	#61422	3' overlap seq	2761	4595
3	ESR1 #74494	#74494	ER-LBD	1	774
4	Insertion site to BamHI	#61422	5' overlap seq	4596	8094

Table 2: Primer pairs for PCR reaction. Annealing temperatures were calculated using Benchling.

#	Primer	Annealing Temp °C	nmol	µL water for 100 µM	Template DNA
1	Insertion-BamI FWD	73.2	16.74	167.4	dCAS9-VP64_GFP #61422
2	Insertion-BamI REV	75	33.47	334.7	dCAS9-VP64_GFP #61422
3	Age1-Insertion FWD	73.1	30.87	308.7	dCAS9-VP64_GFP #61422
4	Age1-Insertion REV	75	29.78	297.8	dCAS9-VP64_GFP #61422
5	ERBD FWD	75	30.27	302.7	pBLO1811_arC9_noNLS_human #74494
6	ERBD REV	72.9	27.23	272.3	pBLO1811_arC9_noNLS_human #74494

Gibson assembly of ESR1 into dCas9-VP64 finalizes the construct:

The dCas9-VP64 host was first validated and then digested at the AgeI and BamHI sites (Figure 3). The outputs of this process were the validation of the ordered dCas9-VP64 plasmid and the isolation of DNA between AgeI and BamHI for ESR1 insertion.

Initial PCR attempts to amplify the desired fragments encountered challenges, notably with the BamHI site, which can be attributed to insufficient elongation time for the larger 3kbp product. Adjustments were made to the annealing temperature and PCR protocol, resulting in successful amplification at the AgeI site, but gel imaging (Figure 2) found that the ERBD insertion and the BamHI amplifications were unsuccessful. For BamHI, the incorrect DNA template was used. For the ERBD, we theorize that the failed reaction is due to improper annealing times and/or temperatures. Ongoing reflection on protocol modifications and next steps will be required to ensure reproducibility and technical rigor.

After successful amplification, Gibson cloning would ligate the three regions and the dCas9-VP64 backbone, creating darC9-VP64.

DOI: 10.36838/IJHSR820.20

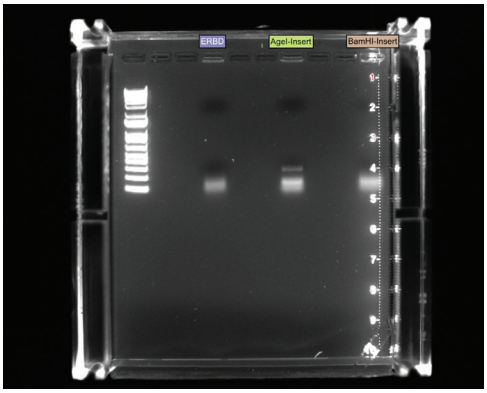


Figure 2: Image of amplification gel cast with ladder on left. Column 1 has ER-LBD primers (5/6 on Table 2), Column 2 has Age1-insert primers (3/4), and Column 3 has BamH1-insert primers (1/2). Column 2 bands are accurate with amplified sequence length (1835bp).

The final darC9-VP64 construct mapped in Figure 4 is a product of the combination of the dCas9-VP64 and arCas9 plasmids. It features a dCas9 with an insertion of the ESR1 sequence at the 231 site.[7] SV40 Nuclear Localization Signals (NLS) flank the darC9 sequence, allowing nuclear entry of the molecule. A GFP is attached, which allows for live cell imaging to track the activity of the complex. The plasmid contains multiple transcriptional boosters, including WPRE and a CMV enhancer. The EF-1α promoter is directly upstream of the ORF containing the darC9-VP64 molecule, while the SV40 promoter comes after.

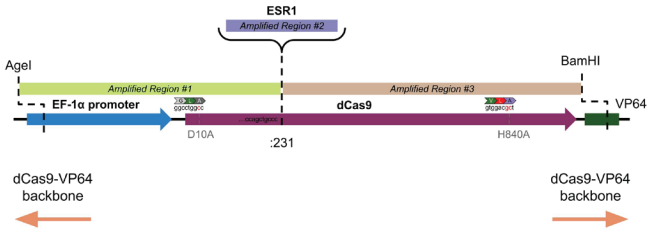


Figure 3: Linear map of darC9-VP64 from the AgeI to BamHI restriction sites, which were used for digestion, amplification, and ESR1 insertion. 3 regions were amplified: from AgeI to the 231 ESR1 insertion site; ESR1; and 231 to BamHI. Annotated are the D10A and H840A mutations that deactivate the RuvC and HNH domains on the Cas9.

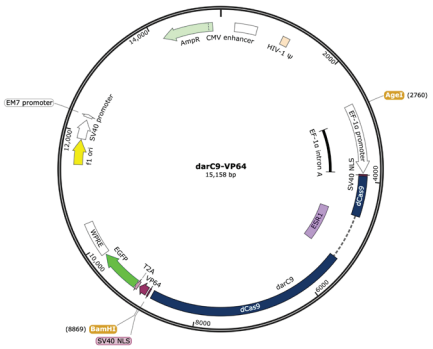


Figure 4: Circular map of darC9-VP64 with selected features and restriction sites.

■ Conclusion

The creation of a darC9-VP64 platform through Gibson cloning and the successful amplification of one of the three fragments demonstrates the plausibility of constructing a dead allosterically regulated Cas9-VP64 tool that could reversibly upregulate a gene after 4-OHT activation.

We attempted to actualize this tool, but we realized that amplification of the ESR1 domain is especially difficult due to GC-rich sections. We successfully amplified the AgeI to ESR1 insertion site section, but could not validate the ESR1 and downstream sections. We theorize that this is due to low primer annealing rates, as a result of imprecise temperatures. Further optimization is required for amplification of the unsuccessful regions. We believe that this can be done through more precise PCR thermal cycles.

Inducible transcriptional activation would be engineered using a dCas9-VP64 fusion gated by an ER-LBD. After sgRNA base pairing, the dCas9-VP64 complex would recognize the target promoter and position itself at the locus. In the absence of the ligand (4-OHT), ESR1 would constrain the complex in an OFF conformation. Administration of 4-OHT would cause an allosteric transformation that permits recruitment of transcription factors. This would enable reversible, time-specific, and tunable control of gene expression.

A functioning darC9-VP64 platform has high potential in therapeutics and gene editing due to its precision. It combines the allostery of arCas9, which has high pertinence for *in vivo* applications requiring incredibly precise timing,[7] with the reversibility and less destructive nature of regulatory CRISPR activation. Compared to other approaches to allostery, such as intein splicing, darC9-VP64 offers greater reversibility, as intein-mediated splicing is a permanent, one-way reaction once triggered. This makes darC9-VP64 better suited for applications requiring variable, temporally precise gene regulation, such as in developmental biology or therapeutics. This system is especially well-suited for applications requiring polygenetic manipulation, as it is especially effective for the coordinated and repeatable regulation of multiple genes.

Future work involving darC9-VP64 will focus on optimizing PCR conditions for the ESR1 and flanking fragments to complete the Gibson assembly and validate the full darC9-VP64 construct through cell assays.

■ Acknowledgments

I would like to thank Dr. Yvon Woappi and the entire Synthetic Regeneration and Systems Physiology lab at Columbia University for their support in mentoring me throughout this research project and throughout the writing of the manuscript.

■ References

1. Ma, Y.; Zhang, L.; Huang, X. Genome Modification by CRISPR/Cas9. *FEBS J.* **2014**, *281* (23), 5186–5193. https://doi.org/10.1111/febs.13110.
2. Jinek, M.; Chylinski, K.; Fonfara, I.; Hauer, M.; Doudna, J.; Charpentier, E. A Programmable Dual-RNA-Guided DNA Endonuclease in Adaptive Bacterial Immunity. *Science* **2012**, *337*, 816–821. https://doi.org/10.1126/science.1225829.
3. Larson, M. H.; Gilbert, L. A.; Wang, X.; Lim, W. A.; Weissman, J. S.; Qi, L. S. CRISPR Interference (CRISPRi) for Sequence-Specific Control of Gene Expression. *Nat. Protoc.* **2013**, *8* (11), 2180–2196. https://doi.org/10.1038/nprot.2013.132.

4. Clouse, G. *CRISPR Activators: A Comparison Between dCas9-VP64, SAM, SunTag, VPR, and More*! https://blog.addgene.org/crispr-activators-dcas9-vp64-sam-suntag-vpr (accessed 2025-07-17).
5. Tanenbaum, M. E.; Gilbert, L. A.; Qi, L. S.; Weissman, J. S.; Vale, R. D. A Protein Tagging System for Signal Amplification in Gene Expression and Fluorescence Imaging. *Cell* **2014**, *159* (3), 635–646. https://doi.org/10.1016/j.cell.2014.09.039.
6. Davis, K. M.; Pattanayak, V.; Thompson, D. B.; Zuris, J. A.; Liu, D. R. Small Molecule-Triggered Cas9 Protein with Improved Genome-Editing Specificity. *Nat. Chem. Biol.* **2015**, *11* (5), 316–318. https://doi.org/10.1038/nchembio.1793.
7. Oakes, B. L.; Nadler, D. C.; Flamholz, A.; Fellmann, C.; Staahl, B. T.; Doudna, J. A.; Savage, D. F. Profiling of Engineering Hotspots Identifies an Allosteric CRISPR-Cas9 Switch. *Nat. Biotechnol.* **2016**, *34* (6), 646–651. https://doi.org/10.1038/nbt.3528.
8. Khehra, N.; Padda, I. S.; Zubair, M. Polymerase Chain Reaction (PCR). In *StatPearls*; StatPearls Publishing: Treasure Island (FL), 2026.
9. Avilan, L. Assembling Multiple Fragments: The Gibson Assembly. *Methods Mol. Biol.* **2023**, *2633*, 45–53. https://doi.org/10.1007/978-1-0716-3004-4_4.
10. Khajanchi, N.; Saha, K. Controlling CRISPR with Small Molecule Regulation for Somatic Cell Genome Editing. *Mol. Ther.* **2022**, *30* (1), 17–31. https://doi.org/10.1016/j.ymthe.2021.06.014.
11. Zetche, B.; Volz, S. E.; Zhang, F. A Split Cas9 Architecture for Inducible Genome Editing and Transcription Modulation. *Nat. Biotechnol.* **2015**, *33* (2), 139–142. https://doi.org/10.1038/nbt.3149.
12. Tang, W.; Hu, J. H.; Liu, D. R. Aptazyme-Embedded Guide RNAs Enable Ligand-Responsive Genome Editing and Transcriptional Activation. *Nat. Commun.* **2017**, *8* (1), 15939. https://doi.org/10.1038/ncomms15939.
13. Zhang, L.; Zhang, X.; Qiu, L.; Mao, S.; Sheng, J.; Chen, L.; Khan, U.; Upputuri, P. K.; Zakharov, Y. N.; Coughlan, M. F.; Perelman, L. T. Near-Infrared Light Activatable Chemically Induced CRISPR System. *Light Sci. Appl.* **2025**, *14*, 229. https://doi.org/10.1038/s41377-025-01917-8.
14. Kundert, K.; Lucas, J. E.; Watters, K. E.; Fellmann, C.; Ng, A. H.; Heineike, B. M.; Fitzsimmons, C. M.; Oakes, B. L.; Qu, J.; Prasad, N.; Rosenberg, O. S.; Savage, D. F.; El-Samad, H.; Doudna, J. A.; Kortemme, T. Controlling CRISPR-Cas9 with Ligand-Activated and Ligand-Deactivated sgRNAs. *Nat. Commun.* **2019**, *10* (1), 2127. https://doi.org/10.1038/s41467-019-09985-2.
15. Shao, J.; Wang, M.; Yu, G.; Zhu, S.; Yu, Y.; Heng, B. C.; Wu, J.; Ye, H. Synthetic Far-Red Light-Mediated CRISPR-dCas9 Device for Inducing Functional Neuronal Differentiation. *Proc. Natl. Acad. Sci. U. S. A.* **2018**, *115* (29), E6722–E6730. https://doi.org/10.1073/pnas.1802448115.
16. Qi, L. S.; Larson, M. H.; Gilbert, L. A.; Doudna, J. A.; Weissman, J. S.; Arkin, A. P.; Lim, W. A. Repurposing CRISPR as an RNA-Guided Platform for Sequence-Specific Control of Gene Expression. *Cell* **2013**, *152* (5), 1173–1183. https://doi.org/10.1016/j.cell.2013.02.022.
17. Zuo, Z.; Liu, J. Allosteric Regulation of CRISPR-Cas9 for DNA Targeting and Cleavage. *Curr. Opin. Struct. Biol.* **2020**, *62*, 166–174. https://doi.org/10.1016/j.sbi.2020.01.013.
18. Liu, K. I.; Ramli, M. N. B.; Woo, C. W. A.; Wang, Y.; Zhao, T.; Zhang, X.; Yim, G. R. D.; Chong, B. Y.; Gowher, A.; Chua, M. Z. H.; Jung, J.; Lee, J. H. J.; Tan, M. H. A Chemical-Inducible CRISPR–Cas9 System for Rapid Control of Genome Editing. *Nat. Chem. Biol.* **2016**, *12* (11), 980–987. https://doi.org/10.1038/nchembio.2179.
19. Bendixen, L.; Jensen, T. I.; Bak, R. O. CRISPR-Cas-Mediated Transcriptional Modulation: The Therapeutic Promises of CRISPRa and CRISPRi. *Mol. Ther. J. Am. Soc. Gene Ther.* **2023**, *31* (7), 1920–1937. https://doi.org/10.1016/j.ymthe.2023.03.024.
20. Strezoska, Ž.; Dickerson, S. M.; Maksimova, E.; Chou, E.; Gross, M. M.; Hemphill, K.; Hardcastle, T.; Perkett, M.; Stombaugh, J.; Miller, G. W.; Anderson, E. M.; Vermeulen, A.; Smith, A. van B. CRISPR-Mediated Transcriptional Activation with Synthetic Guide RNA. *J. Biotechnol.* **2020**, *319*, 25–35. https://doi.org/10.1016/j.jbiotec.2020.05.005.
21. Yeo, N. C.; Chavez, A.; Lance-Byrne, A.; Chan, Y.; Menn, D.; Milanova, D.; Kuo, C.-C.; Guo, X.; Sharma, S.; Tung, A.; Cecchi, R. J.; Tuttle, M.; Pradhan, S.; Lim, E. T.; Davidsohn, N.; Ebrahimkhani, M. R.; Collins, J. J.; Lewis, N. E.; Kiani, S.; Church, G. M. An Enhanced CRISPR Repressor for Targeted Mammalian Gene Regulation. *Nat. Methods* **2018**, *15* (8), 611–616. https://doi.org/10.1038/s41592-018-0048-5.
22. Alerasool, N.; Segal, D.; Lee, H.; Taipale, M. An Efficient KRAB Domain for CRISPRi Applications in Human Cells. *Nat. Methods* **2020**, *17* (11), 1093–1096. https://doi.org/10.1038/s41592-020-0966-x.
23. Qian, J.; Liu, S. X. CRISPR/dCas9-Tet1-Mediated DNA Methylation Editing. *Bio-Protoc.* **2024**, *14* (8), e4976. https://doi.org/10.21769/BioProtoc.4976.
24. McCarty, N. S.; Graham, A. E.; Studená, L.; Ledesma-Amaro, R. Multiplexed CRISPR Technologies for Gene Editing and Transcriptional Regulation. *Nat. Commun.* **2020**, *11* (1), 1281. https://doi.org/10.1038/s41467-020-15053-x.
25. Chardon, F. M.; McDiarmid, T. A.; Page, N. F.; Daza, R. M.; Martin, B. K.; Domcke, S.; Regalado, S. G.; Lalanne, J.-B.; Calderon, D.; Li, X.; Starita, L. M.; Sanders, S. J.; Ahituv, N.; Shendure, J. Multiplex, Single-Cell CRISPRa Screening for Cell Type Specific Regulatory Elements. *Nat. Commun.* **2024**, *15* (1), 8209. https://doi.org/10.1038/s41467-024-52490-4.
26. Hilton, I. B.; D'Ippolito, A. M.; Vockley, C. M.; Thakore, P. I.; Crawford, G. E.; Reddy, T. E.; Gersbach, C. A. Epigenome Editing by a CRISPR-Cas9-Based Acetyltransferase Activates Genes from Promoters and Enhancers. *Nat. Biotechnol.* **2015**, *33* (5), 510–517. https://doi.org/10.1038/nbt.3199.
27. Konermann, S.; Brigham, M. D.; Trevino, A. E.; Joung, J.; Abudayyeh, O. O.; Barcena, C.; Hsu, P. D.; Habib, N.; Gootenberg, J. S.; Nishimasu, H.; Nureki, O.; Zhang, F. Genome-Scale Transcriptional Activation by an Engineered CRISPR-Cas9 Complex. *Nature* **2015**, *517* (7536), 583–588. https://doi.org/10.1038/nature14136.
28. Zhou, H.; Liu, J.; Zhou, C.; Gao, N.; Rao, Z.; Li, H.; Hu, X.; Li, C.; Yao, X.; Shen, X.; Sun, Y.; Wei, Y.; Liu, F.; Ying, W.; Zhang, J.; Tang, C.; Zhang, X.; Xu, H.; Shi, L.; Cheng, L.; Huang, P.; Yang, H. *In Vivo* Simultaneous Transcriptional Activation of Multiple Genes in the Brain Using CRISPR–dCas9-Activator Transgenic Mice. *Nat. Neurosci.* **2018**, *21* (3), 440–446. https://doi.org/10.1038/s41593-017-0060-6.

Author

Rohaan Sanyal is a junior at Sacred Heart Preparatory in Atherton, CA. He is hoping to pursue a degree in the biosciences and a career in translational science.

Appendix A

Below are the summary tables for the literature review mentioned in the Methods section. Some important acronyms included in the following tables include sgRNA, standing for

DOI: 10.36838/IJHSR820.20

short guide RNA, used in CRISPR editing to direct nucleases to desired loci; AAV, standing for Adeno-associated viral vector; and KRAB, standing for Krüppel-associated box, a transcription repressor domain.

Table 3: Compilation of approaches for inducible (conditionally activated) Cas9 systems.

Inducible CRISPR system	Activation domain	Delivery	Control	Limitations/drawbacks	Reference(s)
Intein-mediated split	Cas9 halves fused to split inteins, becoming functional after dimerization in the presence of chemical ligands (4-OHT)	Dual recombinant AAV vectors	Full Cas9 expressed after intein-mediated trans-splicing	Binding efficiency of 4-OHT reduced slightly	[10,11]
Rapamycin-mediated split	dCas9-FRB (N) & dCas9-FKBP-VP64 (C) inteins that reassemble in the presence of rapamycin-induced FRB-FKBP dimerization.	Dual recombinant AAV vectors	Rapamycin induces dimerization and allows full Cas9 expression	Relies on rapamycin. Potential off-target effects.	[11, 12, 13]
Near-infrared (NIR) photoactivatable	Cyanine photobleaching activates rapamycin, inducing Cas9 split-dimerization	Dual recombinant AAV vectors	NIR illumination leads to dimerization	Low penetration depth of NIR	[13]
ligRNA+	Aptamer inserted into sgRNA. Ligand stabilizes complex, allowing for functional dCas9	Cas9 and sgRNA plasmid (produces endogenous ligRNA)	ligRNA aptamer binding turns Cas9-sgRNA on/off	High design complexity. Potential cross-reactivity	[14]
Aptazyme-embedded sgRNAs (agRNAs)	Self-cleaving aptazyme fused to sgRNA. Ligand binding triggers cleavage, releasing the steric blockade strand	Plasmid with dCas9 and agRNA cassettes, transfection	Small molecules (theophylline/guanine) cause aptazyme cleavage	High design complexity. Slower kinetics	[12]
Photoactivatable CRISPR	Light-inducible heterodimerization, photocleavable oligonucleotides, etc. (UV, visible, NIR)	Plasmid expression; electroporated *in vivo* (mouse muscle)	Far red light leads to dimerization, allowing endogenous gene activation + multiplex control	Phototoxicity, limited tissue penetration, slow activation, high design complexity, temperature sensitivity	[15,16]
Allosteric intein split	Same as intein split but with estrogen receptor. Conformational change	Dual recombinant AAV vectors	Conformational change allows for reversibility after Cas9 expression	Inconsistent intein splicing	[6]
ER-LBD arCas9	ER-LBD inserted into Cas9. 4-OHT induces conformational change, removing steric blockade	Plasmid transfection	After 4-OHT binding, ER-LBD switches into antagonist conformation, releasing the REC2 gate	Moderate background activity without inducer. Activation limited by ligand presence & domain arrangement	[17,18]

Table 4: Compilation of approaches for dCas9 interference (CRISPRi) systems.

CRISPRi system	Interference domain	Delivery	Control	Limitations/drawbacks	Reference(s)
Foundational steric blockade (Qi et al.)	dCas9:sgRNA complex bound to nontemplate DNA strand, interfering with transcriptional regulation mechanisms	Two plasmids holding dCas9 and sgRNA chimera, complex assembly	aTc-inducible promoter on Cas9 plasmid; sgRNA constitutive	Moderate repression (~46%) in mammalian cells. Reversibility can take hours to days	[3,16,19]
dCas9-KRAB	dCas9:effector domain complex attaches to target DNA strand, repressor domain recruits transcription repressors	Lentiviral transduction of two vectors, dCas9-repressor domain and sgRNA	Constitutive promoters (SFFV, U6 RNA Pol III)	Variable repression (~75%)	[19,20]
KRAB-MeCP2	KRAB + MeCPE2 tail fused to dCas9	Lentiviral dual-vector	Constitutive promoter	Much larger; MeCP2 can have off-target effects	[21]
ZIM3-KRAB (next gen)	Potent ZIM3 + KRAB	Single lentiviral or PiggyBac	Constitutive promoter	Less explored approach, ZIM3 primate restricted	[22]
dCas9-TET1	Fusion of catalytic demethylation domain to dCas9	Transfection, electroporation, or lentiviral transduction	sgRNA constitutive	Imprecise demethylation	[23]

Table 5: Compilation of approaches for dCas9 activation (CRISPRa) systems.

CRISPRa system	Activation domain	Delivery	Control	Limitations/drawbacks	Reference(s)
dCas9 with activation domains	dCas9:sgRNA complex bound to nontemplate DNA strand, improving transcription efficiency with activation domains (VP64 originally)	Plasmid transfection	Constitutive promoters CMV & U6	Weaker activation, multiple gRNAs required. Off-target effects. Specified architecture & domains.	[24,25]
dCas9-p300	p300 HAT (histone acetyltransferase) core domain	Plasmid transfection	Constitutive promoters CMV & U6	Larger fusion protein. Does not work equally at all loci. Specified optimization	[26]
dCas9 with serial activation domains VPR	VP64-p65-Rta (VPR) domain with dCas9. Serial activation domains more effective	Plasmid transfection, lentiviral vectors	Constitutive promoters	Large fusion protein. Variable efficiency. Possible use of multiple gRNAs.	[5]
dCas9 with SunTag	dCas9 fused with SunTag scaffold for multiple effector domains, improving activation 15-20-fold	Plasmid transfection	Weak promoters	Low expression. Potential interference with protein function due to size.	[5]
SAM (Synergistic Activation Mediator)	dCas9 + VP64 + MS2-p65-HSF1 complex, which recruits additional activation domains	Plasmid transfection, lentiviral vectors	Constitutive promoters	Requires co-expression of multiple components. Decrease in activation efficiency during high multiplexing. Variable gene activation	[27]
SPH (SunTag-p65-HSF1)	SunTag scaffold with p65-HSF1 (from SAM system)	PiggyBac transposon, AAV/lentiviral plasmid	Cre recombinase inducible	Difficult delivery due to large size. Variable gene activation.	[28]
Multiplexed CRISPRa	Multiple gRNAs targeting promoters or enhancers	Lentiviral vectors, piggyBac transposon systems, AAV vectors, plasmid transfection	Various	Off-target effects increase from additional gRNAs.	[24,25]

IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

How Do Aluminum Formwork Systems Perform in High-rise Residential Buildings in India with Respect to Construction Efficiency and Reusability Under Real Site Conditions?

Shivam Sharma

Pathways School Gurgaon, Baliawas Village, Gurugram, Haryana, 122101, India; shivamunimax9@gmail.com
Mentor: Andrea Giusti

ABSTRACT: In India, building tall residential buildings often uses aluminum formwork systems because they allow for quick construction and can be reused. However, there is rarely an assessment of how well these systems perform in real construction sites. To fill this gap, the study looks into how efficient these systems are and how reusable they are in actual site conditions. Using firsthand data from an ongoing high-rise project, such as schedules, labor deployment, on-site observations, and repair records, the study measures efficiency by examining how long it takes to complete each floor and the amount of labor used. Reusability is evaluated by looking at how many times the Formwork can be reused, how often repairs are needed, and how much material is wasted. The results of this study provide practical insights and a method for better planning and management of formwork systems. The study concludes that with proper planning, handling, and management on-site, aluminum formwork can lead to faster construction across multiple floors, fewer repairs, and more reuse cycles.

KEYWORDS: High-rise Construction, Aluminum Formwork Systems, Formwork Reusability, Construction Efficiency, Labor Productivity, Floor Cycle Time.

Introduction

In India, high-rise residential construction is always associated with tight schedules, high material demands, repetitive structural layouts, and labor efficiency requirements. In such projects, effective planning and faster construction are crucial for overall project performance. Construction speed in high-rise buildings is strongly influenced by the type of formwork system used.[1] The selection of an appropriate formwork system directly affects labor productivity, material usage, construction cycle time, and surface finish quality. The relationship between formwork systems and project performance has been examined in several engineering studies by Thiyagarajan *et al.* and Hurd.[2,3] Formwork is a temporary system used to support, hold, and shape freshly placed concrete until it gains sufficient strength to carry loads. Formwork accounts for a significant portion of the total cost and time in reinforced concrete structures, especially in wall-slab systems. Studies on high-rise construction also show that repetitive floor layouts improve the efficiency and reusability of formwork systems.[4] Aluminum formwork systems are designed for repetitive high-rise construction as they are modular and panel-based. These systems use prefabricated aluminum panels to cast slabs and walls in a largely monolithic manner, allowing construction cycles to be repeated floor by floor. Previous studies comparing aluminum and conventional formwork systems indicate that the lower weight and higher dimensional accuracy of aluminum panels enable faster erection and de-shuttering, which reduces construction time and labor dependence.[5] These studies also suggest that aluminum formwork systems can achieve high reuse cycles when managed properly. However, several studies have clearly suggested that the actual performance of aluminum formwork systems depends strongly on site-related factors such as handling practices, storage conditions, workforce skill levels, and planning quality.[1,6] In practice, aluminum formwork performance is often judged based on supplier recommendations or past experience rather than direct measurement on construction sites. This creates a research gap, as project-based information on construction speed, labor productivity, reuse cycles, repair frequency, and material losses is rarely documented for Indian high-rise projects. Without empirical evidence, engineers and project managers rely mainly on assumptions that may not reflect actual site performance. Prior studies of formwork systems have similarly indicated the need for empirical, site-specific data so that assumed performance benefits can be validated.[1,4] Therefore, a site-based evaluation is required to support technical decision-making and improve formwork management practices in the Indian construction context. This study addresses this gap by evaluating the performance of aluminum formwork systems at a real high-rise residential construction project in India. The central research question is: Do aluminum formwork systems, when properly planned, handled, and managed, give better and improved construction efficiency and higher reuse performance across successive floors in real site conditions? The study specifically examines construction efficiency using floor cycle time and labor deployment as indicators, and material reusability using achieved reuse cycles, repair frequency, and material wastage. This study addresses this gap by carrying out a site-based evaluation of aluminum formwork systems used in a high-rise residential construction project in India. The study focuses on construction efficiency,

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.21

measured using indicators such as floor cycle time and labor requirement, and material reusability, measured using achieved reuse cycles, repair frequency, and material wastage.

Material and Methods

This study follows an empirical, site-based research approach. The analysis is based mainly on primary data collected directly from a real high-rise residential construction project in India. Secondary literature is used only to support background understanding, selection of performance indicators, technical calculations, and standard engineering practices.

Study Area and Project Description:

The selected project for this case study research is Sobha International City, located in Sector 109, Gurugram, near the Dwarka Expressway. Refer to Figures 1 and 2. The project is developed by SOBHA Limited, a reputed, well-established real estate developer known for large-scale residential projects and quality construction practices.

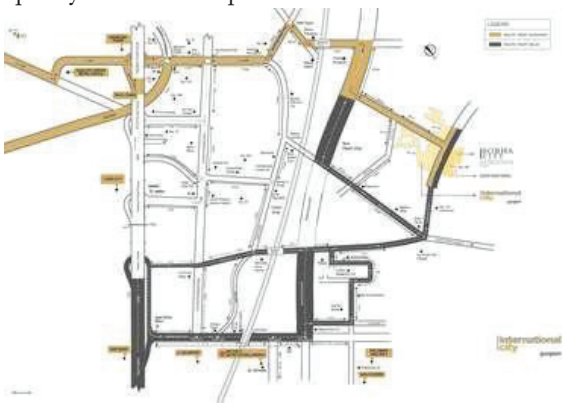


Figure 1: Project location map: Sobha International City, Sector 109, Gurugram.

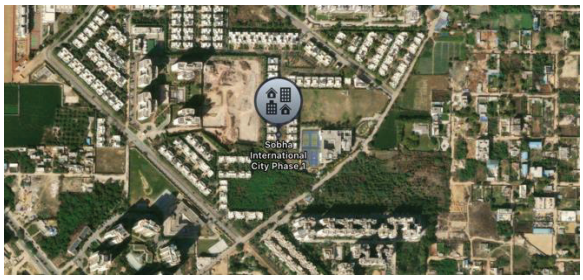


Figure 2: Satellite image: Sobha International City.

The project consists of high-rise residential towers with repetitive floor layouts and reinforced concrete wall-slab structural systems, as shown in Figure 3. Aluminum formwork is used for the wall and slab construction on multiple floors. The repetitive structural configuration and organized construction layout make this project suitable for evaluating construction efficiency and material reusability of aluminum formwork systems. The location near the Dwarka Expressway provides good accessibility for material movement, labor mobilization, and equipment operation, which supports continuous construction activities.

Refer to Table 1 for project information such as location, building type, number of floors, and typical floor area.

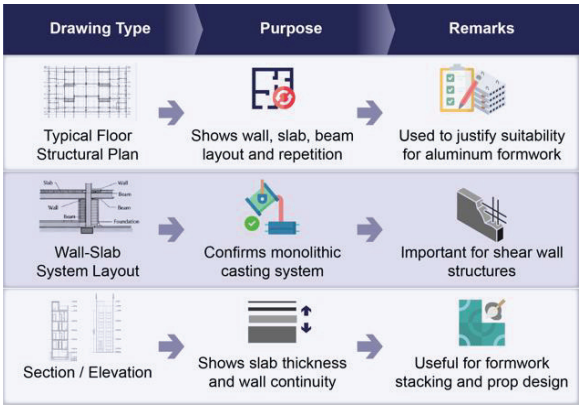


Figure 3: Analysis of structural drawings justifying aluminum formwork system suitability.

Table 1: Project information gathered from the site engineer at Unimax International, part of the Max Group.

Parameter	Details
Project Name	Sobha International City
Location	Sector 109, Gurugram, Haryana, India
Developer	SOBHA Limited
Project Type	High-rise residential development
Structural System	Reinforced concrete wall-slab system
Formwork System Used	Aluminum formwork system
Number of Towers (Studied)	2
Number of Floors per Tower	G+20 to G+30
Typical Floor Area	3000–3500m^2 per floor
Construction Method	Monolithic wall and slab casting
Site Layout	Organized layout with repetitive floor configuration
Connectivity	Located near Dwarka Expressway

Data Sources and Data Collection:

Data were collected through direct site-based observation and review of construction records. Laboratory testing was not involved in this study. Observations were carried out during ongoing construction activities to understand actual site conditions, including formwork erection, concreting, stripping, handling, and storage of aluminum formwork panels.

Primary Data Sources:

Primary data were collected from the following sources:

Site observations during the erection, concreting, and de-shuttering of formwork; construction schedules showing actual floor cycle durations; labor deployment records for typical floors; reuse and repair records maintained by the formwork team; interviews and informal discussions with site engineers and supervisors; using multiple data sources helped improve the reliability, accuracy, and consistency of the collected information. Appendices referenced: Appendix A: Sample of site observation sheets and Appendix B: Interview questionnaire format.

Performance Indicators:

The following parameters were selected for evaluation using commonly adopted indicators from previous formwork studies:

Construction Efficiency, Floor cycle time (days per floor), Labor deployment per floor, Reusability, Number of reuse cycles achieved, Repair frequency, Material wastage

These indicators are widely used in studies related to formwork performance.

Procedure

- Data collection was carried out floor by floor during regular site visits. The sequence of activities followed during data collection included:
- Observing aluminum formwork erection and de-shuttering processes across multiple floors
- Recording actual floor cycle times from construction schedules
- Collecting labor deployment data from site records
- Reviewing reuse and repair information with the formwork team
- Cross-checking collected data with site engineers to ensure accuracy
- To ensure consistency across different floors and reduce errors in reporting, this procedure is followed.

Engineering Design Checks:

Standard engineering checks related to fresh concrete pressure, load on props, panel deflection, and safe de-shuttering were considered as part of normal construction practice to ensure safe and repeatable use of aluminum formwork systems on site. These checks were to maintain safety and quality controls during construction.

Detailed formulas, assumptions, and sample calculations based on relevant codes and standards are presented separately in Appendix C.

Data Analysis Method:

The collected data, which refer to Figure 4, were analyzed using simple comparative and trend-based methods: Planned and actual floor cycle times were compared to study construction efficiency; labor deployment trends were reviewed across repeated floors; reuse cycles and repair records were analyzed to identify patterns related to panel type and handling practices.

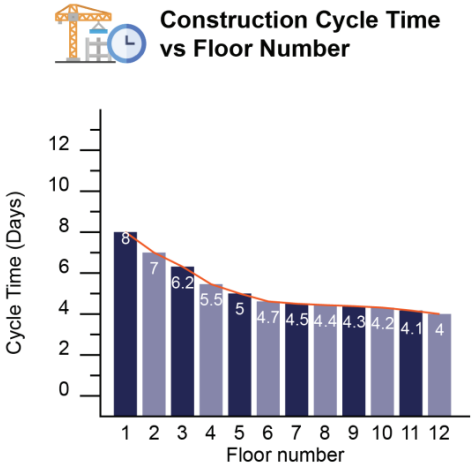


Figure 4: Decreasing floor-wise construction cycle over progressive floors.

Data Limitations and Adjustments:

During the Initial stages of construction, complete records for some floors were not available, and repair information for a few panels was not clearly updated. These issues were taken care of through discussions with the site supervisor and by observing additional floors. In cases where exact values were not available, average values were used. Such adjustments are commonly adopted in practical construction site studies.

Results and Discussion

Construction Efficiency:

The aluminum formwork system requires a floor cycle time of 8 to 9 days per floor in the studied project. Previous studies on conventional formwork systems report typical floor cycle times in the range of 12 to 15 days per floor, which are used here only for reference from the literature.

Table 2: Construction efficiency benchmark.

Parameter	Present Study (Aluminum Formwork)	Literature Benchmark (Conventional Formwork)
Observed Floor Cycle Time	8–9 days per floor	12–15 days per floor
Data Source	Site construction schedules, Sobha International City, Gurugram	Thiyagarajan *et al.* (2017)[1]; Supriyono *et al.* (2025)[4]
Construction Type	Repetitive high-rise residential, R.C. wall-slab system	Conventional timber/plywood formwork, multi-story R.C. structures
Remarks	Cycle stabilized after 3–4 repetitions following initial learning period	Typically longer due to sequential assembly and lower panel reusability

This comparison clearly shows a difference in construction time when aluminum formwork is used under repetitive conditions. During the initial floors of the project, the cycle time was slightly higher due to a combination of factors. These included the initial learning period as workers became familiar with the aluminum formwork system, as well as the greater structural complexity typically associated with lower floors. Lower floors commonly contained more complex reinforcement layouts, higher fresh concrete pressure due to larger concrete head, a higher density of openings such as doors and windows, and more detailed corner and junction connections. These conditions increased the time required for accurate panel setting, alignment, and support, which contributes to longer cycle times independent of worker experience. However, after 3 to 4 repetitions, the structural layouts became more uniform, and workers better developed familiarity with the aluminum formwork system. As a result, the cycle time became stable and consistent, indicating improved workflow and coordination. Labor deployment ranged between 70 and 90 workers per floor. The proportion of skilled labor was approximately 25 to 30 percent, while the remaining workforce mainly consisted of semi-skilled labor. For this research study, skilled labor refers to workers who have previous experience in aluminum formwork operations, including carpenters and formwork supervisors who can independently interpret construction drawings, align panels correctly, and operate lifting equipment. Semi-skilled labor refers to workers with general construction site experience who are assigned repetitive tasks such as carrying panels, prop placement, and routine cleaning, under the direct supervision of skilled workers. These labor categories are consistent

DOI: 10.36838/IJHSR820.21

with standard labor classifications generally followed on Indian construction sites.

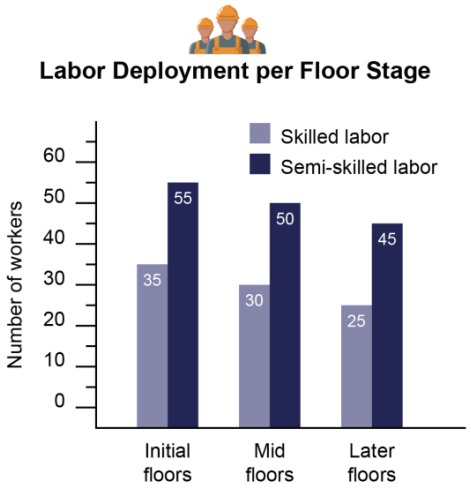


Figure 5: Labor deployment for aluminum formwork operations is decreasing per floor stage over time.

As the number of repetitions increased, workers became more familiar with the aluminum formwork system, which led to a gradual improvement in labor productivity. Overall, these results show that aluminum formwork supports shorter and more stable construction cycles in repetitive high-rise construction. Figure 5 shows the number of workers engaged per floor during aluminum formwork operations.

As shown in Figure 6, the floor cycle time reduced and stabilized after repeated use of the aluminum formwork system, reflecting a learning curve effect.

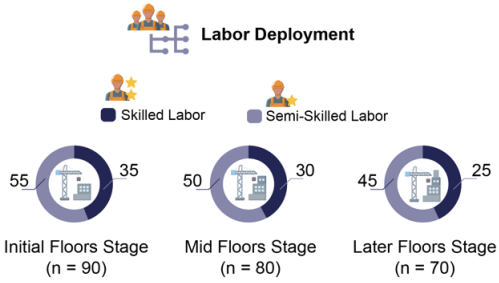


Figure 6: Learning curve effects in aluminum formwork: stabilization of cycle time and reduction in labor deployment across repeated floors.

Reusability:

Aluminum formwork systems offer high reuse potential (refer to Figure 7 and Figure 13) and long service life when compared to timber and plywood formwork systems, which can only be reused a limited number of times. Properly maintained aluminum panels can be reused across multiple construction cycles. According to Thiyagarajan *et al.* (2017), aluminum formwork panels are capable of achieving up to 200 to 300 reuse cycles, reducing the need for frequent replacement of formwork materials. High repetition reduces the total quantity of formwork required over the project duration and lowers long-term construction costs. Aluminum panels with good durability also maintain dimensional accuracy and shape over repeated uses, which helps in achieving consistent concrete quality and reduces rework. In this project, standard wall and deck panels have shown high rates of repetition with minimal damage, supporting the benefits in terms of durability reported in earlier studies.

In terms of long-term material performance, aluminum panels maintain dimensional accuracy and surface condition across repeated cycles when proper maintenance is followed. According to Thiyagarajan *et al.* (2017),[1] aluminum formwork panel is capable of achieving 200 to 300 reuse cycles under appropriate handling and storage conditions. High repetition reduces the total quantity of formwork required over the project duration and lowers long-term construction costs. The high reuse potential observed in the present project is consistent with these published values and confirms that the durability of aluminum formwork is closely tied to site management practices rather than material limitations alone.

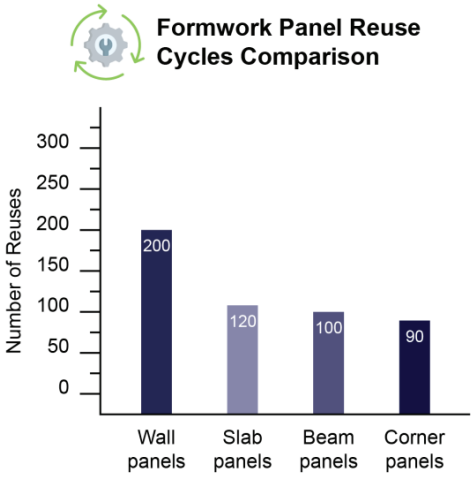


Figure 7: Higher reuse cycles achieved by different aluminum panel types.

These findings indicate that aluminum formwork systems can achieve high reuse levels when proper inspection and maintenance practices are followed on site. Figure 7 shows how the number of reuse cycles changes among different aluminum panel types. Figure 8 shows how often repairs are required and the types of damage observed during repeated use. Site observations gave a practical insight into the on-site performance of aluminum formwork systems.

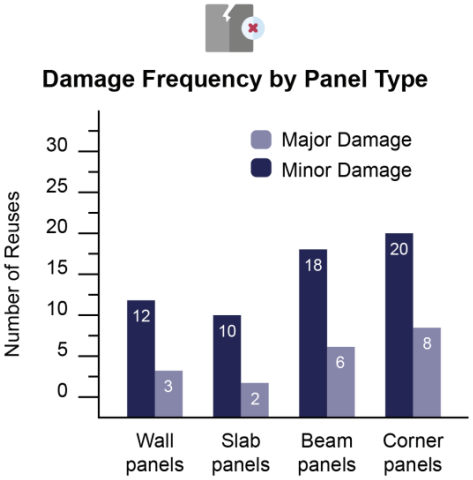


Figure 8: Damage frequency by panel type: corner and beam panels exhibit the highest susceptibility to both minor and major damage.

Most panel damage was observed during the stripping and shifting stages. In several cases, rough handling during the removal of panels caused bending at the edges and minor surface dents. Panels that were handled carefully showed fewer defects and could be reused for a greater number of cycles.



Figure 9: Proper stacking of panels for reducing contact damage and surface wear.

Storage practices also had a significant impact on panel condition. Panels that are stored in designated storage areas with wooden battens placed in between are found in better condition than panels stacked directly on the ground. Proper stacking of panels reduced contact damage and surface wear significantly. These observations show that handling and storage practices play a critical role in maintaining the overall quality of formwork and achieving higher reuse rates.



Figure 10: The damaged panel during de-shuttering was caused by improper handling.

Figure 9 shows the organized storage of aluminum panels using wooden battens to prevent damage. Figure 10 presents common panel damage caused by improper handling during de-shuttering.

Sustainability, Cost Efficiency, and Environmental Impact:

Sustainability performance is an important consideration in modern high-rise construction, particularly in projects where repetitive use, material efficiency, and environmental impact are major concerns. Aluminum formwork systems are increasingly adopted not only for construction speed and surface quality, but also for their reusability, recyclability, durability, and long-term environmental benefits. This section discusses the sustainability, cost efficiency, and environmental impact of aluminum formwork systems based on site observations and supported by published research.

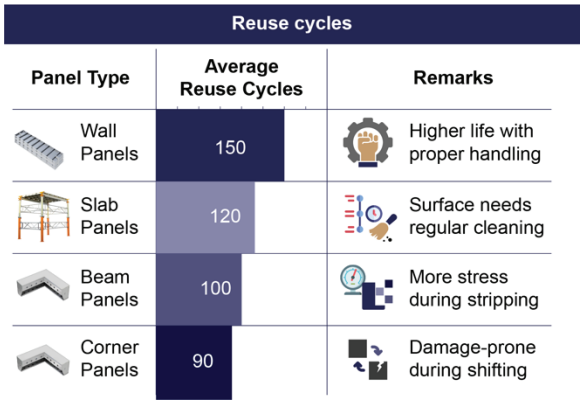


Figure 11: Wall panels exhibit maximum longevity among formwork components.

Aluminum can be recycled multiple times (refer to Figure 11) without loss of quality or strength, making aluminum formwork suitable for sustainable construction practices. According to the International Aluminum Institute, more than 90% of aluminum used in construction is recyclable, and producing aluminum from recycled material requires only about 5% of the energy needed to produce aluminum from raw bauxite.[7] This high recyclability means that:

Scrap aluminum obtained from damaged or retired panels can be recycled. End-of-life formwork systems do not become landfill waste. Material value is recovered at the end of service life. In the present project, damaged panels and scrap aluminum were either returned to the supplier or sold to authorized recyclers, supporting circular economy practices (refer to Figure 12).



Figure 12: Scrap aluminum storage and recycling arrangement on site supporting circular economy practices.

Carbon Emissions in Manufacturing and Life-Cycle Impact:

The environmental impact of aluminum formwork systems depends mainly on the energy source used during aluminum production. According to the International Aluminum Institute, aluminum produced using hydropower has a carbon

DOI: 10.36838/IJHSR820.21

footprint of approximately 10 tCO_2e per ton, compared to approximately 30 tCO_2e per ton for coal-based production.[8] In India, electricity generation is largely dependent on coal and thermal power, which increases emissions during aluminum production. However, life-cycle assessment studies show that although initial emissions may be higher, the high reuse rate of aluminum formwork offsets this impact over time. Life-cycle studies on reusable formwork materials report that aluminum formwork systems have a lower overall environmental impact when evaluated over multiple reuse cycles compared to disposable timber formwork. This shows that aluminum formwork systems are environmentally favorable when assessed over their full-service life rather than only at the manufacturing stage, as shown in Figure 13.

Parameter	Reusable Aluminum Formwork	Conventional Wooden Formwork
Lifespan	>250 times (Over 6 years)	Up to 20–40 time (less than 1 year)
Reusability and Durability	High reusability, resistant to wear and tear	Prone to wear, swelling and degradation
Maintenance	Low maintenance, simple cleaning and repair	Frequent repairs, prone to damage in handling
Initial Cost	High, with return on investment over time	Low, but higher overall across multiple projects
Overall Life Cycle Cost	Lower in long term due to higher lifespan & reusability	Higher due to frequent replacement and repairs

Figure 13: Aluminum's extensive reusability offsets initial manufacturing emissions, lowering overall life-cycle environmental impacts.

De-Shuttering Efficiency and Time Savings:

Aluminum formwork systems allow safe and early de-shuttering because of their accurate alignment and efficient support systems. According to IS 456:2000, formwork may be removed once the concrete has gained sufficient early strength to carry applied loads safely.[9]

Aluminum formwork systems enable:

1. Faster stripping of wall and slab panels
2. Reduced damage to concrete edges and corners
3. Shorter construction cycle times
4. Early de-shuttering helps reduce overall project duration
5. Equipment and crane running time
6. Indirect energy consumption on site

In the present project, faster de-shuttering contributed to achieving stable floor cycles of 8–9 days per floor, as discussed in Section 3. Refer to Figure 14 for a stripping time comparison of aluminum and conventional formwork systems.

Aluminum formwork provides high dimensional accuracy and a smooth surface finish, which improves the quality of cast concrete. According to M.K. Hurd,[3] rigid and well-aligned formwork systems reduce honeycombing, concrete leakage, and surface defects. Improved surface finish often eliminates the need for plastering and reduces remedial works. In the present project, only skim coating finishing was required, and plastering was largely eliminated. This resulted in: Lower consumption of cement, sand, and finishing materials; reduced labor requirement; lower energy use related to rework. Refer to Figure 15 for the concrete surface finish after de-shuttering

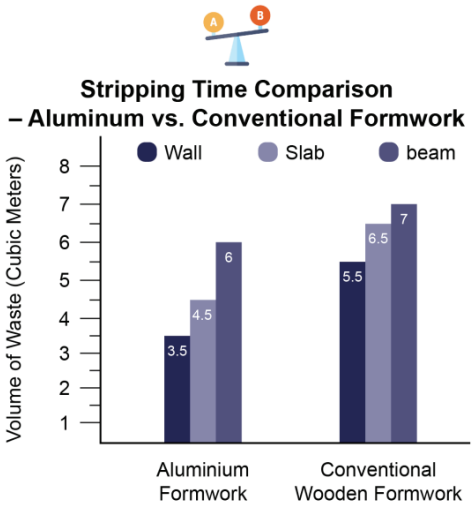


Figure 14: Aluminum formwork reduces stripping time across all components compared to conventional wooden formwork.



Figure 15: High-quality concrete surface finish achieved directly after de-shuttering.

Reduction in Construction Waste:

Traditional timber and plywood formwork systems generate high levels of construction waste. Refer to Figure 16 for a comparison of waste generation: aluminum vs conventional formwork, because traditional formwork is reused fewer times and disposed of after limited cycles. According to the Central Pollution Control Board (CPCB),[10] construction and demolition waste is a major contributor to landfill volumes in urban India. Reusable aluminum formwork systems reduce timber consumption, plywood disposal, and on-site waste generation. In the present project, timber and plywood waste were minimal due to the extensive use of aluminum formwork.

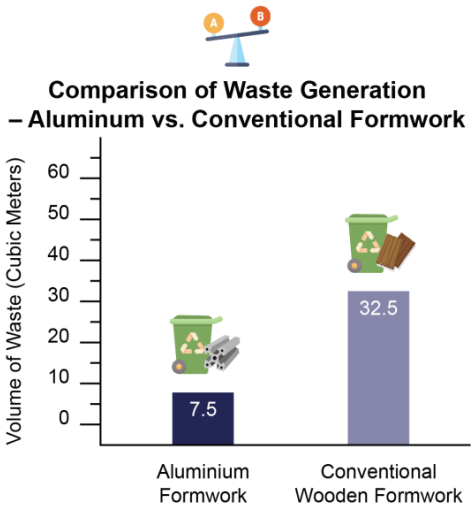


Figure 16: Aluminum formwork drastically reduces construction waste compared to wood.

Although aluminum formwork systems require a higher initial investment, their sustainability characteristics lead to long-term cost efficiency. According to Comparative studies published in IJRESM and IJARET, high reuse reduces cost per use; faster construction cycles reduce overhead costs; and reduced finishing work lowers labor and material expenses.[4,11,12] As a result, aluminum formwork systems become economically favorable in large projects with repetitive layouts. Refer to Figure 17 for a cost comparison between aluminum & conventional formwork systems.

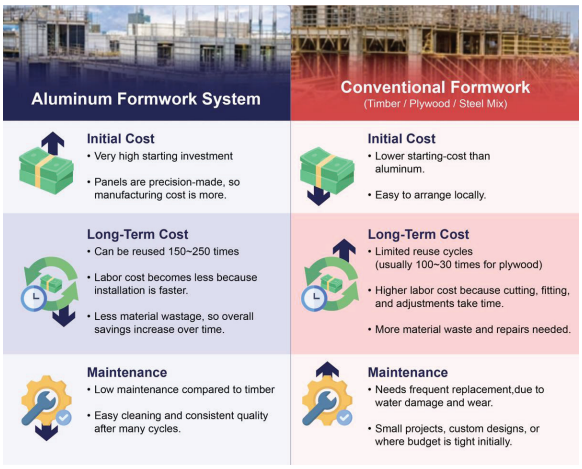


Figure 17: Comparison of Aluminum and Conventional Formwork Systems.

Contribution to Sustainable Development Goals (SDGs):

Aluminum formwork systems support different "United Nations Sustainable Development Goals". They support SDG 9 by promoting the use of advanced construction techniques, SDG 11 by improving housing delivery speed and reducing construction waste, and SDG 12 by supporting reuse, recycling, and circular construction practices. According to the United Nations Environment Program, material reuse and recycling play an important role in achieving sustainable urban development.[13]

Overall Sustainability Assessment:

From an environmental and practical perspective, aluminum formwork systems give several advantages, such as high reusability, full recyclability, reduced construction waste, improved concrete quality, faster construction cycles, and lower long-term cost. Although aluminum manufacturing may involve carbon emissions depending on the energy source, the overall life-cycle performance of aluminum formwork remains favorable when evaluated over multiple reuse cycles.

Effect of Planning and Construction Sequencing:

In this project, proper planning and construction sequencing played an important role in achieving shorter floor cycle times. After the initial learning period, the construction cycle stabilized at 8–9 days per floor, indicating that repetitive sequencing improved productivity and coordination. According to Thiyagarajan *et al.* (2017), aluminum formwork systems always perform better in projects with repetitive floor layouts and fixed construction cycles, as stable sequencing reduces delays and improves coordination between labor and activities.[1] Similarly, M.K. Hurd explains that well-planned formwork operations reduce idle time, ensure continuity in concreting, and minimize rework. In this study, floors completed after the sequencing process became more stable, showed fewer interruptions, and had more predictable cycle times. This supports the hypothesis that improved planning and sequencing contribute directly to better construction efficiency. Refer to Figure 18 for a typical construction sequence flowchart using aluminum formwork.

These findings of this project directly support the study's hypothesis that better planning and construction sequencing directly improve construction efficiency, as evidenced by the stable 8-9 day floor cycle achieved following sequence stabilization in the present project.

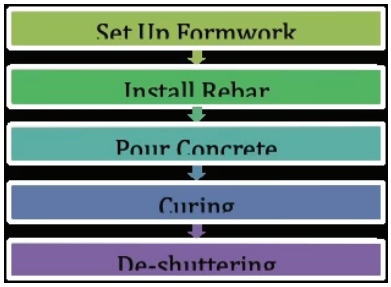


Figure 18: Construction sequence flowchart for improved construction efficiency.

Influence of Labor Skill and Training:

Labor skill level had a clear effect on construction speed and panel condition. During the initial floors, higher cycle times were observed due to unfamiliarity with aluminum formwork systems. After 3–4 repetitions, labor productivity improved and cycle times reduced, refer to Figure 19.

Formwork productivity studies reported in IJARET show that labor efficiency is directly connected with workers' experience in system formwork, due to learning curve effects in repetitive construction activities. Thiyagarajan *et al.* (2017) also reported that trained formwork teams achieve faster erection

DOI: 10.36838/IJHSR820.21

and de-shuttering with fewer alignment errors and lower damage rates.

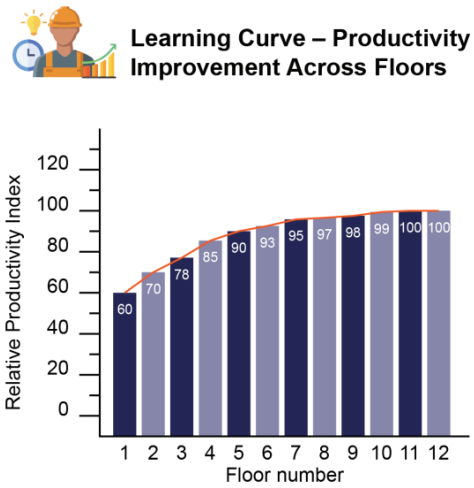


Figure 19: Repeated floor cycles significantly improve relative labor productivity.

In the present project, floors constructed after worker training showed:

1. Faster erection and stripping
2. Reduced dependence on skilled carpenters
3. Lower repair frequency

This indicates that labor training and experience are critical for improving construction efficiency and reducing damage to formwork panels, which supports the hypothesis as well.

This confirms the hypothesis that site management practices, specifically workforce training and skill development, are critical determinants of both construction speed and material performance and directly address the study's research question regarding efficiency under real site conditions.

Effect of Handling Practices during Stripping and Shifting:

The handling methods used for de-shuttering and panel shifting significantly affected the frequency of repairs and damage. In most of the observed cases, it is shown that bending and deformation of edges occurred due to improper stripping or careless handling during shifting. According to M.K. Hurd, poor stripping techniques can cause deformation of panels and reduce their service life.[3] ACI 347 also highlights the importance of careful removal and proper lifting methods to avoid damage to both formwork panels and concrete edges.[14] In the present project, refer to Figure 20; it is shown that panels that were stripped carefully and lifted using correct procedures showed fewer repairs, better dimensional stability, and higher reuse potential. This confirms that handling practices directly influence panel durability and reuse performance.

This finding directly supports the hypothesis by demonstrating that handling quality is a controllable site factor that governs reuse performance outcomes, independent of the inherent material properties of aluminum formwork panels.



Figure 20: Correct handling and stripping procedures directly enhance panel reuse performance.

Influence of Storage Conditions on Reuse Performance:

Storage conditions had a direct impact on reuse performance. Panels that are stored in designated yards with wooden battens and proper stacking showed less deformation, surface damage, and alignment issues.

According to Thiyagarajan *et al.* (2017), improper stacking and exposure to moisture increase the risk of bending and surface damage.[1] The International Aluminum Institute also states that protection from physical damage improves the service life of aluminum products. In the present study, panels stored with wooden separators and without direct ground contact achieved higher reuse cycles and required fewer repairs. Figure 21 shows correct stacking and storage of aluminum formwork panels, which shows that storage methods play an important role in controlling repair frequency and extending panel life.

This reinforces the study's central argument that when aluminum formwork systems are properly managed, including through adequate storage discipline, they demonstrably deliver higher reuse performance, as proposed in the hypothesis.



Figure 21: Correct stacking of aluminum formwork panels using protective wooden separators directly extends panel reuse life.

Relationship between Site Practices and Reusability:

The reuse cycles recorded in this project, ranging from 200 to 250 repetitions, are consistent with values reported in earlier studies. Thiyagarajan *et al.* (2017) reported that, by following proper handling, storage, and maintenance practices,

aluminum panels can achieve up to 200 reuse cycles.[1] However, panels exposed to frequent impact, poor cleaning, or improper stacking showed earlier damage and required more frequent repairs. This indicates that reusability is influenced not only by material properties but also by site management practices.

Impact on Repair Frequency and Maintenance Cost:

Repair frequency in the project was directly related to handling methods and the type of panels used. Beam and corner panels required earlier maintenance due to higher stress concentration, complex geometry, and repeated handling. Studies reported in IJRESM have clearly indicated that frequent minor repairs can increase maintenance costs if preventive measures are not adopted.[4,11,12] In the present project, regular inspection and minor repairs after every 10–15 cycles helped prevent major damage and avoid early panel replacement. This shows that timely maintenance improves panel durability and reduces long-term cost.

Overall, these findings confirm that repair frequency should not be considered a constant material property, but rather a factor that changes management practices and that timely maintenance protocols are very important for achieving the reuse potential that the hypothesis anticipated.

Sustainability Implications and Environmental Performance:

Beyond construction efficiency and reusability, the data collected from a real site in this study also offer insights into the broader sustainability performance of aluminum formwork systems. The high reuse cycles observed in this project directly support the environmental and cost-efficiency claims reported in prior research, and are discussed here in the context of life-cycle impact and sustainable construction practice.

Overall Interpretation of Results:

The discussion shows that although aluminum formwork systems have high technical potential, their actual performance depends mainly on: Planning and sequencing quality, Labor skill and training, handling during stripping and shifting, Storage discipline, and regular inspection and maintenance. These findings are consistent with earlier formwork studies. The results support the hypothesis that improved site practices lead to better construction efficiency and higher reuse performance. Aluminum formwork should therefore be considered not only as a material system, but as a management-controlled construction process, where final efficiency, reuse, and cost performance are determined by site practices.

■ Conclusion

This study evaluated the performance of aluminum formwork using data collected from a high-rise residential project at Sobha International City, Gurugram. The results show that aluminum formwork delivers faster construction cycles, high reuse, improved surface finish, reduced waste generation, and long-term cost efficiency as compared to other formwork systems. The engineering calculations are used properly in this project to ensure safe operation and acceptable construction quality. At the same time, the findings also highlight the role of aluminum formwork in supporting sustainable construction practices and relevant SDG goals. The study demonstrates the importance of site-based evaluation and provides practical guidance for planning and managing aluminum formwork in future high-rise residential projects in India.

Recommendations for Construction Managers:

1. Based on the findings of this study, the following practical recommendations are made for construction managers planning to deploy aluminum formwork systems on high-rise residential projects:

2. Labor training should be conducted before starting formwork operations, with at least 3–4 supervised practice floors recommended to reduce the learning-curve effect on cycle time. Training of workers should specifically cover panel alignment, prop setting, and safe stripping techniques.

3. Storage protocols should require that all aluminum panels be stacked in designated storage yards on wooden battens, with a minimum of 50 mm separation between panels, and protected from direct ground contact and exposure to impact during shifting.

4. Regular inspections should be scheduled after every 10–15 reuse cycles for standard wall and deck panels, while beam and corner panels should be checked after every 8 to 10 cycles, which are more susceptible to deformation due to geometric complexity and handling stress.

5. De-shuttering operations should be carried out using designated lifting hooks and mechanical support equipment. Site instructions should clearly discourage manual dragging or dropping of panels during stripping.

Recommendations for Future Research:

This study is based on data from a single high-rise residential project in Gurugram, Haryana. Future research studies should broaden the scope by evaluating the aluminum formwork systems across several project sites in different cities of our country and climatic zones, including those regions that have high humidity, extreme temperature variation, or coastal exposure, which may affect panel degradation rates. Comparative studies between aluminum formwork and steel formwork systems under similar site conditions would also provide useful quantitative benchmarks. Additionally, a full life-cycle cost analysis that includes initial procurement, maintenance, repair, transportation, and recycling value at the end of service life gives strength to the economic case for aluminum formwork adoption in the Indian construction sector.

■ Acknowledgments

I would like to thank my research mentor, Andrea Giusti, for his guidance throughout the research process, and Unimax International, part of the Max Group, for providing access to site data and technical insights that contributed to the development of this research paper.

DOI: 10.36838/IJHSR820.21

References

1. Thiyagarajan, R.; V.Panneerselvam; Nagamani, K. ALUMINIUM FORMWORK SYSTEM USING IN HIGHRISE BUILDINGS CONSTRUCTION. *Int. J. Adv. Res. Eng. Technol. IJARET* **2017**, *8* (6), 29–41.
2. Kannan, M. R.; Santhi, M. H. Automated Constructability Rating Framework for Concrete Formwork Systems Using Building Information Modeling. *Asian J. Civ. Eng.* **2018**, *19* (4), 387–413. https://doi.org/10.1007/s42107-018-0026-3.
3. m. k. hurd. *Formwork for Concrete*; the American Concrete Institute, 1979.
4. Supriyono, A.; Putra, I. N. D.; Trigunarsyah, B. Comparative Analysis of Formwork Systems: Cost Efficiency and Time Management in Construction Projects. *Adv. Sustain. Sci. Eng. Technol.* **2025**, *7*, 0250204. https://doi.org/10.26877/asset.v7i2.1320.
5. Auliyah Shabirah, S. Comparative Analysis of Formwork Systems: Cost Efficiency and Time Management in Construction Projects. *Adv. Sustain. Sci. Eng. Technol*. **2025**. https://doi.org/10.26877/asset.v7i2.1320.
6. Abreu, M.; Casado, A. Aluminum Formwork System: Loss and Productivity. *Built Environ. Proj. Asset Manag*. **2019**, ahead-of-print. https://doi.org/10.1108/BEPAM-04-2018-0070.
7. Stacey, Michael. Demolition, Reuse and Recycling. In *Aluminum Recyclability and Recycling: Towards Sustainable Cities*; Cwningen Press: Nottingham, 2015; pp 30, 54–55.
8. International Aluminum Institute. *Aluminum Carbon Footprint FAQs*. International Aluminum Institute. https://international-aluminum.org/landing/aluminum-carbon-footprint-faqs/ (accessed 2026-04-28).
9. *IS 456:2000 — Plain and Reinforced Concrete: Code of Practice*; 4th Revision; Bureau of Indian Standards IS 456:2000; Bureau of Indian Standards: New Delhi, 2000.
10. Ministry of Environment, Forest and Climate Change. *Environment (Construction and Demolition) Waste Management Rules, 2025*; 2025.
11. Pawar, Sandip P.; Atterde, P. M. Comparative Analysis of Formwork in Multistory Building. *International Journal of Research in Engineering and Technology* **2014**, *3* (Special Issue 09), 22–24.
12. Kendre, C.; Bhamare, T. Comparison between Aluminum and Conventional Formwork Based on Cost and Time. *Int. J. Res. Eng. Sci. Manag.* **2024**, *7* (2), 71–73.
13. United Nations Environment Program. *Buildings and Climate Change: Status, Challenges and Opportunities*; United Nations Environment Program: Nairobi, 2007. https://wedocs.unep.org/handle/20.500.11822/7783.
14. ACI Committee 347. *Guide to Formwork for Concrete*; ACI 347-04; American Concrete Institute, 2004.

Authors

Shivam Sharma is an International Baccalaureate Diploma Program student at Pathways School, Gurgaon, India. His academic interests include structural civil engineering and sustainable construction systems. His research focuses on improving efficiency and reusability in high-rise residential formwork systems.

Appendix A

Project Information

Project Name: Sobha International City

Location: Sector 109, Gurugram

Date of Observation: _____

Floor Number: _____

Observer Name: _____

1. Activity Details

Observation Item	Details
Construction Stage	□ Formwork Erection □ Reinforcement □ Concreting □ De-shuttering
Panel Type Observed	□ Wall □ Slab □ Beam □ Corner
Work Area	Room / Core / Slab / Wall Line
Start Time / End Time	_____ / _____

2. Labor Deployment

Category	Number of Workers
Skilled Labor	_____
Semi-Skilled Labor	_____
Helpers	_____
Total Labor	_____

3. Formwork Condition Check

Check Item	Yes / No	Remarks
Panels aligned properly	□ Yes □ No	_____
Wall straightness checked	□ Yes □ No	_____
De-shuttering oil applied	□ Yes □ No	_____
Leakage control (timber support)	□ Yes □ No	_____
Proper stacking of panels	□ Yes □ No	_____

4. Damage / Repair Observation

Panel Type	Damage Observed	Action Taken
Wall Panel	_____	_____
Slab Panel	_____	_____
Beam Panel	_____	_____

5. Safety & Handling Practices

Proper lifting method used: □ Yes □ No

Rough handling noticed: □ Yes □ No

Storage yard condition: □ Good □ Average □ Poor

6. Cycle Time Notes

Planned activity for the day: _____

Actual progress observed: _____

Delay reason (if any): _____

7. Additional Remarks

_____

Appendix B - Site Engineer/Supervisor Interview Questionnaire

Project: Aluminum Formwork Performance Study – High-Rise Residential Construction

Purpose: To collect practical insights from site engineers and supervisors regarding construction efficiency, reusability, handling practices, and on-site challenges.

General Information

1. Name of Engineer / Supervisor:

_____

2. Role on Site:

_____

3. Years of Experience in Formwork Systems:

_____

4. Project Name / Tower:

_____

5. Date of Interview:

_____

Section A – Construction Efficiency

6. What is the typical floor cycle time observed using aluminum formwork?

7. What are the main reasons for delay during initial floors?

8. How does labor productivity change after repeated cycles?

9. What challenges are faced during erection and de-shuttering?

Section B – Reusability and Maintenance

10. What is the average reuse cycle achieved for wall, slab, and beam panels?

11. Which panel types require frequent repair or maintenance?

12. What are the common causes of panel damage on site?

13. How often are inspections carried out for aluminum panels?

Section C – Handling and Storage Practices

14. What procedures are followed during stripping and shifting of panels?

15. How aluminum panels stored after de-shuttering?

16. Are wooden battens or protective measures used during stacking?

Response from the interview

1) Project Information - Sobha International City

• Project Name: Sobha International City
• City / State: Gurugram, Haryana, India
• Type of Building: High-rise residential development
• Number of Floors: Typically, G+20 to G+30 (varies by individual tower)
• Typical Floor Area: Approximately 3,000–3,500 m^2 per typical floor

2) Structural Drawings

• Typical Floor Structural Drawings:
o Reinforced concrete shear walls (internal and external)
o Reinforced concrete slabs (typical thickness 150–200 mm)
o Integrated/concealed beams where applicable

• Structural System Adopted:
o RCC shear wall-slab structural system

• Drawings Used for Formwork Planning:
o Structural Drawings
o Architectural Drawings
o MEP Drawings

3) Aluminum Formwork Layout & Panel Details

• Aluminum Formwork Layout Drawings:
o Wall panel layouts (internal and external)
o Slab soffit and deck layouts
o Beam side and soffit panel layouts
o Prop and deck support layouts

• Panel Schedule:
o Panel types: Wall panels, deck panels, beam panels, corner panels, kicker panels, Closing panels
o Standard Wall panel size: Width 600 mm; height 2400mm
o Approximate quantity: 2,000–2,500 panels per tower (indicative)

• Approximate Panel Weights:
o Standard Wall panels: Width 600mm, Height 2400mm-27 kg
o Standard Deck panels: Width 600mm, Height 1200mm-13.5 kg

4) Conventional Formwork Data

• Typical Floor Cycle Time:
o Aluminum formwork: 8–9 days per floor
o Conventional formwork: 12–15 days per floor

• Activity Sequence:
o Layout → Reinforcement fixing → Formwork erection → MEP Work → Supporting MS

• Change in Cycle Time with Progress:
o Initial floors showed longer cycle times due to learning effects
o Cycle time stabilized after 3–4 repetitions

5) Labor Data

• Number of Workers per Floor: 70–90 workers
• Skill Level Distribution:
o Skilled labor: approximately 25–30%
o Semi-skilled labor: approximately 70–75%

• Labor Requirement Across Cycles:
o Labor productivity improved progressively with repetition
o Dependence on skilled carpentry labor reduced after initial cycles

6) Reuse Cycle Records

• Number of Reuse Cycles Achieved: 200–250 cycles per aluminum panel
• Reuse Limits Followed on Site:
o Supplier-recommended Repetition limits
o Periodic inspection and quality control checks

• Variation in Reuse Between Panel Types:
o Standard panels exhibited higher repetition potential
o Beam and corner panels required earlier maintenance

7) Repair & Damage Records

• Frequency of Repairs: Minor repairs typically required after every 10–15 repetitions
• Common Types of Damage:
o Edge bending and deformation

• Main Causes of Damage:
o Improper handling during stripping and shifting
o Concrete leakage due to inadequate sealing
o Improper stacking and storage practices

8) Material Wastage Data

DOI: 10.36838/IJHSR820.21

• Quantity of Scrap Aluminum Generated:- It depends on the Next project.
• Handling and Storage Practices:
o Designated storage yards
o Panels stacked with wooden battens for separation and protection.

• Disposal / Recycling Method:
o Scrap aluminum sold to authorized recyclers or returned to vendors

9) Quality & Finish Observations
• Surface Finish Achieved: Smooth and uniform concrete surface
• Plastering / Remedial Work Requirement:
o Plastering eliminated
o Only putty required

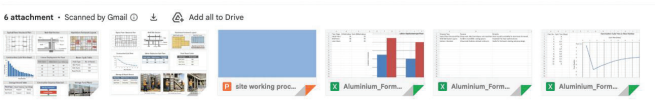
• Quality-Related Site Observations:
o Improved dimensional accuracy and verticality
o Reduced rework and construction defects

10) Site Photographs

11) Cost Data (Indicative)
• Approximate Formwork Cost: ₹10,500–11,500 per m^2 (initial investment)
• Repair / Replacement Cost: Approximately 3–5% of total formwork cost over project lifecycle
• Labor Cost Difference: 20–25% reduction compared to conventional formwork

12) Supervisor Insights
• Ease of Use: High after initial training period
• Speed of Construction: Significantly improved due to repetitive and standardized operations
• Training Requirements: Approximately 7–10 days for workforce familiarization
• Practical Challenges:
o High initial capital investment
o Requirement for strict sequencing and coordination
o Changes in drawings cannot be made once the shell plan is approved

13) Other Construction Techniques / Conventional Formwork



Appendix C

Fresh Concrete Pressure on Formwork

$$P = \gamma \times h$$

Where:
P = pressure on formwork (kN/m^2)
γ = unit weight of fresh concrete (≈ 24 kN/m^3)
h = height of concrete (m)
This formula calculates wall pressure to prevent bending, leakage, and failure of aluminum wall panels. Reference: M.K. Hurd Concrete Formwork, ACI 347 Guide to Formwork for Concrete.

Pouring Rate and Temperature Effect

$$P = C \times R \times T$$

Explanation: This calculation helps control the increase in pressure caused by fast pouring and temperature changes, so that panels are not suddenly overloaded. Reference: ACI 347R Formwork for Concrete.

2.5.3. Load on Props:

$$W = A \times t \times \gamma$$

Explanation: This calculation is used to check the load coming on props so that slab collapse and excessive settlement during concreting can be avoided. Reference: IS 14687 Falsework and Formwork Code.

2.5.4. Spacing of Props:

Explanation: Proper spacing in props is important to give support to the slab during concreting. If the spacing is more, extra deflection may occur, and if it is too little, then unnecessary material and cost are added. Reference: IS 14687 - Falsework and Formwork Code, M.K. Hurd[3] Concrete Formwork.

2.5.5. Bending Stress in Aluminum Panels:

$$\sigma = M / Z$$

Explanation: This check is important to make sure that aluminum panels do not bend permanently during concreting. By controlling the bending stress, the shape of the panels is maintained, and the same panels can be reused many times. Reference: Eurocode 9 Design of Aluminum Structures.

2.5.6. Deflection Limit:

$$\delta = (5wL\text{^}4)/384El$$

Explanation: Controls surface finish quality and prevents uneven walls and honeycombing.
Reference: IS 456,[9] M.K. Hurd.[3]

2.5.7. Safe De-Shuttering Time:

$$ft \geq 0.7fck$$

Explanation: Ensures concrete gains sufficient strength before stripping to avoid cracks and collapse. Reference: IS 456:2000

2.5.8. Productivity Calculation :

$$productivity = (floor\ area)/(cycle\ time \times labour\ floor\ area)$$

Explanation: Measures construction speed and labor efficiency per floor cycle. Reference: IJRESM, IJARET.

IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

Challenges and Possible Treatments for HER2-associated Breast Cancer

TIEN, YU-HSI

Taipei Wego Private Senior High School, No. 50, Zhuhai Rd, Zhongxin Village, Beitou District, Taipei City, 11254, Taiwan (ROC); cecilia20081020@gmail.com

ABSTRACT: Globally, breast cancer is the most common type of cancer in women. Human epidermal growth factor receptor 2 (HER2) overexpression is a significant cause of breast cancer. Breast cancer is categorized into different subtypes according to the level of HER2 expression. This categorization ensures the most efficient treatment is given to every patient. HER2-low breast cancer (BC) is a novel subtype that has been discovered. With fewer HER2 proteins detected, many treatments and former anti-drug conjugates (ADCs) are not useful in curing these patients. Nevertheless, several clinical trials have shown that trastuzumab deruxtecan (T-DXd), a type of ADC, is helpful in the treatment of HER2-low BC cases. This review aims to distinguish different subtypes of breast cancer, treatments available, the mechanism and efficacy of trastuzumab deruxtecan, and identify current challenges and future potential in breast cancer treatment.

KEYWORDS: Medical and Health Sciences, Pharmaceutical Sciences, HER2 Breast Cancer, HER2-Low, Trastuzumab Deruxtecan.

Introduction

In 2025, an estimated 319,750 women were diagnosed with breast cancer. [1] This is not only happening in the US, but it is a worldwide issue. Breast cancer is common in women between the ages of 50 and 69, who have inherited cancer mutations, poor health status like obesity, habits like smoking and drinking alcohol, and experience of radiation exposure or hormone replacement therapy.[2] These will all cause breast cancer to develop when cancer cells abnormally multiply and eventually create a tumor in the breast.

Human epidermal growth factor receptor 2 (HER2), estrogen receptor (ER), and progesterone receptor (PR) are proteins that are deeply associated with breast cancer. HER2, estrogens, and progesterone all help cancerous cells grow. The status of the receptors' expression creates different subtypes of breast cancer that require different treatments. ER-positive (ER+) tumors have estrogen receptors; PR-positive (PR+) tumors have progesterone receptors; and HR-positive (HR+) tumors have both. Another way to categorize breast cancer types is by measuring the amount of HER2 protein. They can be categorized into HER2-positive and HER2-negative. HER2-negative includes HER2-low and HER2-zero.

Anti-HER2 targeting agents have been used in HER2-positive cases and have proved to be helpful. 2 Common anti-HER2 drugs include trastuzumab, pertuzumab, trastuzumab deruxtecan, and trastuzumab.[3] HER2-negative tumors receive endocrine therapies, chemotherapies, or other treatments, depending on the positivity of ER and PR. HER2-low breast cancer is a refined category within HER2-negative. About 60% of HER2-negative tumors are categorized as HER2-low, with a small amount of HER2 but not entirely without expression.[4]

Despite traditional treatments, discoveries in new treatments shed light on the field of research. HER2-targeting drugs were never thought to be effective in HER2-low cases with low HER2 expression. Recent studies have proved that trastuzumab deruxtecan (T-DXd), a type of antibody drug conjugate (ADC), has shown remarkable progress in stopping breast cancer progression and extending patients' survival span. Adaptation of T-DXd can provide a more efficient HER2-low BC treatment, benefiting those who used to be categorized as HER2-negative. This review will introduce different types of HER2-associated breast cancer, current treatment for each of them, and challenges and future potential of breast cancer therapies. This review aims to provide a clear overview of evolving clinical data and treatments surrounding HER2-associated breast cancer, challenges, and possible solutions.

Discussion

1. Types of HER2-associated Breast Cancer: HER2-positive, HER2-low, HER2-zero:

Breast cancer is a heterogeneous disease, or a disease that has more than one cause. In breast cancer tumors, HER2 expression can be the cause of the abnormal growth of cancerous cells. Figure 1 illustrates the HER2 signaling pathway. HER2 receptors do not have a known ligand. Instead, it is activated through dimerization with other receptors from the family, HER3, HER4, or itself, HER2. Overexpression of the HER2 receptor proteins switches on permanent dimerization and activation. Several types of amino acid mutations were also found to be responsible for triggering dimerization of HER-family receptors. [6] Even when growth factors are absent, HER2 dimers keep signaling. Thus, cells bypass normal checkpoints easily, resist apoptosis, and proliferate excessively. This uncontrollable cell growth leads to the formation of tumors. According to the status of HER2 expression, breast cancer can be divided into

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR820.22

different types: HER2-positive, HER2-low, and HER2-zero. Immunohistochemistry (IHC) and in situ hybridization (ISH) are two methods that are used to distinguish the distinct subtypes. This process is important because patients of different types respond positively to different treatments. The next section will introduce the methods to identify them and the current treatment available.

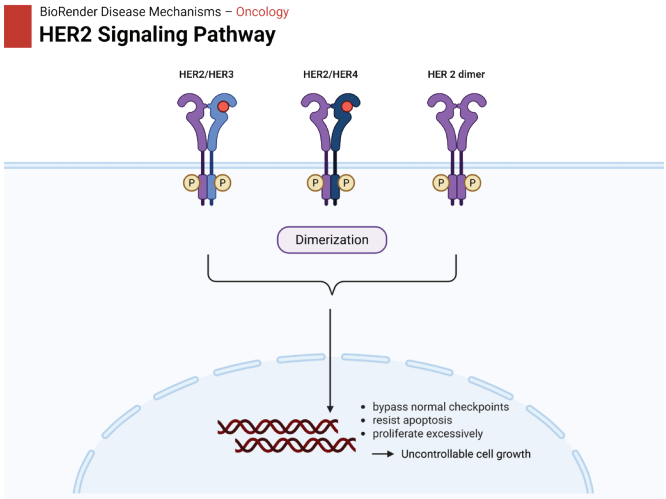


Figure 1: HER2 Signaling Pathway in Breast Cancer. When the HER2 protein undergoes homodimerization (HER2 dimerizes with HER2) or heterodimerization (HER2 dimerizes with other receptor proteins in the family), it triggers uncontrollable cell growth by allowing the tumor cell to bypass normal checkpoints, resist apoptosis, and proliferate excessively. (Created in https://BioRender.com)

1.1. Methods Used to Distinguish HER2 Breast Cancer Subtypes:

To distinguish different types of HER2 breast cancer, IHC analysis is usually conducted first, and ISH is later used to further identify the amount of expression. IHC staining is a staining method that detects and visualizes specific proteins we are looking for. The IHC score ranges from 0, 1+, 2+, to 3+ depending on the amount of visible staining in the cells. Figure 2 shows how HER2 expression can be visualized on IHC staining results. A higher score represents more expression.

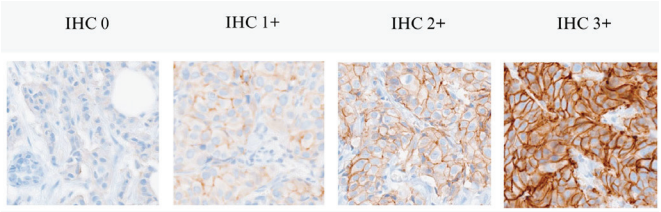


Figure 2: IHC Score With Visualized Results.[5] Tumor cells with higher HER2 expression are presented with a darker color and higher density on the IHC staining results.

In breast cancer tumors, the amount of staining of the HER2 protein is the evaluation standard of IHC scores. IHC 3+ has the most obvious HER2 protein expression, and IHC 0 has the least.[6] Additionally, tumors with an IHC score of 2+ have to take an ISH test to confirm their classification as either HER2-positive or HER2-low. Unlike IHC, ISH aims to test the amplification of the HER2 gene. ISH positive identifies the overexpression of the HER2 gene. Tumors with IHC 3+ or IHC 2+ with an ISH positive are considered HER2-positive tumors. Tumors with IHC 1+ or IHC 2+ with an ISH negative are considered HER2-low. Lastly, tumors with IHC 0 are diagnosed as HER2-zero. Table 1 shows the clearer version that demonstrates the evaluations.

Table 1: HER2 Breast Cancer Subtypes Evaluations.

HER2 BC Subtype	IHC score (0, 1+, 2+, 3+)	ISH result (positive/negative/not required)
HER2-positive	3+	Not required
	2+	Positive
HER2-low	2+	Negative
	1+	Not required
HER2-zero	0	Not required

1.2. Available Treatments for HER2-positive Breast Cancer: Comparison:

Among all breast cancers, 15-20% are HER2-positive tumors.[7] These cases with overexpression and amplification of HER2 have a worse prognosis and a high recurrence rate. Fortunately, anti-HER2 therapies have created new and more advanced treatment options for patients with HER2-positive breast cancer. Currently, trastuzumab and pertuzumab are used in first-line treatment, and trastuzumab emtansine and trastuzumab deruxtecan (T-DXd) in second-line treatment.[3] Trastuzumab is a type of humanized monoclonal antibody that binds to HER2 receptors to down-regulate the pathway. It binds to three specific regions of domain IV of the HER2 extracellular domain.[8] Although the efficacy of trastuzumab treatment is shown to be controversial among different study trials regarding treatment duration and whether the tumor is ligand-dependent or ligand-independent, findings suggest that trastuzumab inhibits the cell cycle of tumor cells through suppressing downstream pathways.[8] For instance, the H0648g phase 3 study significantly proved the effectiveness of trastuzumab. This clinical trial compared the effectiveness of treatment with chemotherapy only and with both trastuzumab and chemotherapy.[9] 469 patients with HER2-positive metastatic breast cancer were involved in the study, where 234 patients received chemotherapy alone, and 235 patients received chemotherapy combined with trastuzumab. The results of the study show that the addition of trastuzumab to chemotherapy contributed to a significantly longer time to disease progression, longer duration of response, a higher rate of objective response, and a lower rate of death. Trastuzumab was approved by the U.S. Food & Drug Administration (FDA) in 1998 for HER2-positive metastatic breast cancer.

Additionally, another study compared the effectiveness and safety of the drugs used in second-line treatment, trastuzumab emtansine (T-DM1) and trastuzumab deruxtecan (T-DXd).[10] 524 patients with HER2-positive tumors were randomly assigned to receive either trastuzumab emtansine or trastuzumab deruxtecan. The result shows that patients treated with trastuzumab deruxtecan have a higher possibility of survival, survival without disease progression, and a higher rate of objective response. Thus, trastuzumab deruxtecan has been proven to be

superior to trastuzumab emtansine in the second-line treatment. This is mainly due to the higher drug-to-antibody ratio of approximately 8 deruxtecan released after the binding of the drug to HER2 receptors. However, trastuzumab deruxtecan is also associated with side effects, especially interstitial lung disease or pneumonitis.

The other group, the HER2-negative tumors, can receive endocrine therapies. HER2-negative treatments depend on the expression of other proteins involved in the cancer. For example, if estrogen is high in the breast cancer tumor, patients can take drugs that lower the estrogen level or block the receptor.

2. HER2-low Breast Cancer and Trastuzumab Deruxtecan:

HER2-low breast cancer is a new category established from the HER2-negative subtype. HER2-low breast cancer tumors have a lower amount of HER2 protein expression compared to HER2-positive tumors; however, they have more than HER2-zero tumors. Before the discovery of trastuzumab deruxtecan, the amount of HER2 protein wasn't enough for other HER2-targeting drugs to detect. Therefore, HER2-low tumors were treated as HER2-negative tumors, mostly receiving endocrine therapies. Currently, HER2-low tumors are defined by IHC 1+ or 2+ with ISH negativity, and about 60% of HER2-negative tumors are now considered HER2-low. The following section will introduce the characteristics of HER2-low tumors and compare T-DXd with other therapies.

HER2-low tumors are more common in HR-positive than in HR-negative tumors compared to HER2-zero.[11] HR-positive tumors are those that overexpress both estrogen and progesterone receptors. This creates the opportunity for a positive response to endocrine therapies targeting the expressed hormones. HER2-low tumors were also proven to have a better prognosis than HER2-zero tumors.[12]

Trastuzumab deruxtecan is an antibody drug conjugate (ADC) that targets the HER2 receptor protein to block growth signaling and damage the DNA, which leads to apoptosis of tumor cells. This paragraph will discuss the function and mechanism of ADC and trastuzumab deruxtecan. When the HER2 gene is mutated, it uncontrollably produces HER2 receptor proteins that regulate cell growth, repair, and apoptosis. ADCs are the "biological missiles," composed of a monoclonal antibody and a potent drug. They selectively bind to target proteins efficiently and accurately, making them a rising star in the research of cancer therapies.[13] Trastuzumab deruxtecan combines trastuzumab, the monoclonal antibody, and deruxtecan, the potent drug. They are chemically linked to each other and are infused into the veins so they can be transported to the proteins they are looking for, in this case, the HER2 receptor. Trastuzumab is like a delivery; it transports the payload, deruxtecan, to the HER2 receptor on the membrane of the tumor cells. Then the complex is ferried into the tumor cells, and the chemical linkage between trastuzumab and deruxtecan breaks. Deruxtecan is a type of topoisomerase I inhibitor. As a weapon carried by trastuzumab, when it is released into the tumor cell, it damages the DNA and leads the cell to apoptosis. Moreover, deruxtecan can easily cross the membrane, entering and attacking neighboring tumor cells regardless of the amount of target protein (HER2) expression. This bystander effect allows it to become one of the most effective drugs in breast cancer therapies and develops its potential to cure patients with HER2-low breast cancer.[2,14,15] Figure 3 demonstrates how T-DXd works on tumor cells expressing the HER2 protein.

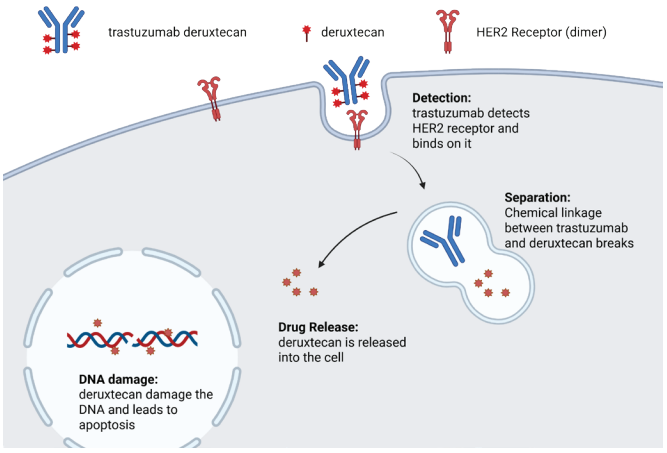


Figure 3: Mechanism of Trastuzumab Deruxtecan. Trastuzumab binds to HER2 receptors and crosses the membrane. After entering the cell, deruxtecan is released to damage DNA, leading to the apoptosis of the tumor cell. (Created in https://BioRender.com)

The DESTINY-Breast04 clinical trial proved the effectiveness of T-DXd in HER2-low breast cancer treatment.[4] 557 patients with HER2-low tumors defined by an IHC 1+ or 2+ with negative results of ISH were involved in this phase III clinical trial.

All of these patients had received one or two chemotherapies before participating in this study. The patients were randomly grouped to compare the T-DXd and standard chemotherapy effectiveness on progression-free survival. The results show that patients treated with T-DXd have significantly longer progression-free survival and overall survival. Thanks to the DB-04 trial, T-DXd became the first approved ADC for use on metastatic HER2-low breast cancer in 2022. This anti-HER2 drug, previously used on HER2-positive tumors, is now developing future potential for ADC treatment in HER2-low and HER2-zero tumors.

3. Challenges and Future Potentials:

New and more advanced drugs and treatments have been developed to target different subtypes of metastatic breast cancer. They establish future potentials based on similar mechanisms, but also face challenges such as the doubtful identification of subtypes, resistance to drugs, and side effects of these drugs. What are the solutions to these issues? What are the future directions of breast cancer treatment studies? This section will introduce the current challenges and future potentials.

The current HER2 status evaluation system was established to identify HER2-positive tumors instead of distinguishing HER2-low tumors from HER2-zero tumors.[2] Since the evaluation of the HER2 breast cancer subtype is based on IHC staining, the results might seem different to different inter-

DOI: 10.36838/IJHSR820.22

preters. The IHC score is based on the visible result of the staining area, which can be inaccurate and doubted. In a recent study, 18 pathologists were asked to read and score 170 breast cancer biopsies. The result shows that the concordance rate of distinguishing IHC 0 from IHC 1+ was 26% among the 18 pathologists. In other words, the diagnosis of HER2-low (IHC 1+) or HER2-zero (IHC 0) can be different based on the pathologists' observations. This is an issue as new ADCs or other drugs are approved to be used in HER2-low tumor treatment. Patients who are inaccurately diagnosed might lose the opportunity to receive such therapies that might be helpful to them. Therefore, a new and improved version of the evaluation of HER2 status should be developed and replace the current one. One solution to this issue is the application of quantitative continuous scoring (QCS). QCS is a learning model that can help with detecting the expression of the HER2 protein by visualizing IHC staining results. Analysis shows that the correlation of QCS results and pathologists' results is high (R=0.993), proving that it can help overcome the conflict with the identification of HER2 expression status.[16]

Although antibody drug conjugates (ADCs) have become drugs that are phenomenal in anticancer research, resistance to them is a challenge to scientists. Since ADCs' mechanism has multiple steps, resistance to the drug can happen at any level.[17] In antigen recognition and expression, the most concerning issue is the status of HER2 expression. Continuous exposure to ADCs can lower the amount of HER2 expression, making the tumor cells with lower HER2 proteins thrive without being killed by the potent drugs.[18] Moreover, ADCs like trastuzumab emtansine (T-DM1) and trastuzumab deruxtecan (T-DXd) are usually used in second-line treatment. Treatment from previous lines might have reduced the expression of HER2, for example, chemotherapy, trastuzumab, and pertuzumab. Resistance might occur if the ADC cannot detect and bind to the receptors. The bystander effect can overcome this issue.[17] ADCs like trastuzumab deruxtecan have a better ability to cross the membrane and enter other tumor cells regardless of the amount of expression. Resistance can also occur after the ADC enters the tumor cells, known as intracellular resistance.

Currently, researchers are committed to finding new and more advanced drugs and treatments for metastatic HER2-associated breast cancer. ADCs like trastuzumab deruxtecan are waiting for approval and are expected to have great potential in increased efficacy and decreased risks, resistance, and side effects.

Conclusion

Breast cancer is one of the most common cancers in the world, especially among women. HER2, estrogens, and progesterone are three factors that can trigger abnormal cell growth. Tumors with overexpression of HER2 are HER2-positive, often treated with anti-HER2 drugs. HER2-low and HER2-zero subtypes were derived from the HER2-negative group, having low or no expression of HER2 protein. These tumors often receive endocrine therapies, depending on the expression of the hormones. However, trastuzumab deruxtecan (T-DXd), an antibody drug conjugate that was previously used in second-line treatment for HER2-positive tumors, was proven to be helpful in HER2-low tumor treatment. The discovery of such drugs set up the future directions for research in breast cancer, but also created some issues that have yet to be resolved. To provide precise medication, the categorization of HER2 breast cancer subtypes using IHC score can be updated. This will allow clearer identification of HER2-low tumors from HER2-negative tumors, allowing more efficient treatments as new drugs like T-DXd are developed. Adaptation of more standardized systems can also improve the accuracy in identification by increasing the concordance of pathologists' observations and evaluations. In response to efficacy and resistance, more research diving into the mechanism of drugs and HER2 dimerization can suggest more helpful treatment combinations. Ultimately, by marrying objective, machine-learning diagnostics with highly tailored combination therapies, the medical community can move beyond general classification systems and transition into a highly individualized, highly adaptable era of cancer care.

Acknowledgments

I would like to thank Dr. Shaye for leading me through the process of writing, and our teaching assistant, Dr. Lauren Tetz, for giving me helpful feedback.

References

1. Siegel, R. L.; Kratzer, T. B.; Giaquinto, A. N.; Sung, H.; Jemal, A. Cancer Statistics, 2025. *CA. Cancer J. Clin.* **2025**, *75* (1), 10–45. https://doi.org/10.3322/caac.21871.
2. Kang, S.; Kim, S.-B. HER2-Low Breast Cancer: Now and in the Future. *Cancer Res. Treat.* **2024**, *56* (3), 700–720. https://doi.org/10.4143/crt.2023.1138.
3. Stanowicka-Grada, M.; Senkus, E. Anti-HER2 Drugs for the Treatment of Advanced HER2 Positive Breast Cancer. *Curr. Treat. Options Oncol.* **2023**, *24* (11), 1633–1650. https://doi.org/10.1007/s11864-023-01137-5.
4. Modi, S.; Jacot, W.; Yamashita, T.; Sohn, J.; Vidal, M.; Tokunaga, E.; Tsurutani, J.; Ueno, N. T.; Prat, A.; Chae, Y. S.; Lee, K. S.; Niikura, N.; Park, Y. H.; Xu, B.; Wang, X.; Gil-Gil, M.; Li, W.; Pierga, J.-Y.; Im, S.-A.; Moore, H. C. F.; Rugo, H. S.; Yerushalmi, R.; Zagouri, F.; Gombos, A.; Kim, S.-B.; Liu, Q.; Luo, T.; Saura, C.; Schmid, P.; Sun, T.; Gambhire, D.; Yung, L.; Wang, Y.; Singh, J.; Vitazka, P.; Meinhardt, G.; Harbeck, N.; Cameron, D. A. Trastuzumab Deruxtecan in Previously Treated HER2-Low Advanced Breast Cancer. *N. Engl. J. Med.* **2022**, *387* (1), 9–20. https://doi.org/10.1056/nejmoa2203690.
5. Bornstein-Quevedo, L.; de Anda-González, J.; Lara-Torres, C. O.; Flores-Gutiérrez, J. P.; Dorantes-Heredia, R.; Bautista-Piña, V.; Zaragoza-Vargas, P.; Alcaraz-Wong, A.; Soto-Sañudo, A. K.; Mendoza-Ramírez, S.; Salamanca-García, M.; Loyola-Rodríguez, G.; Gómez-Macías, G. S.; Murguia-Perez, M.; De Luna-Sánchez, M.; Villalobos-Valencia, R.; Talamantes, E.; Arce-Salinas, C. Navigating HER2-Low Testing in Invasive Breast Cancer: Update Recommendations for Pathologists. *J. Pers. Med.* **2024**, *14* (5), 467. https://doi.org/10.3390/jpm14050467.
6. Wu, Q.; Xu, L. Challenges in HER2-Low Breast Cancer Identification, Detection, and Treatment. Transl. *Breast Cancer Res.* **2024**, *5*, 3–3. https://doi.org/10.21037/tbcr-23-48.

7. Li, Y.-W.; Dai, L.-J.; Wu, X.-R.; Zhao, S.; Xu, Y.-Z.; Jin, X.; Xiao, Y.; Wang, Y.; Lin, C.-J.; Zhou, Y.-F.; Fu, T.; Yang, W.-T.; Li, M.; Lv, H.; Chen, S.; Grigoriadis, A.; Jiang, Y.-Z.; Ma, D.; Shao, Z.-M. Molecular Characterization and Classification of HER2-Positive Breast Cancer Inform Tailored Therapeutic Strategies. *Cancer Res.* **2024**, *84* (21), 3669–3683. https://doi.org/10.1158/0008-5472.CAN-23-4066.
8. Maadi, H.; Soheilifar, M. H.; Choi, W.-S.; Moshtaghian, A.; Wang, Z. Trastuzumab Mechanism of Action; 20 Years of Research to Unravel a Dilemma. *Cancers* **2021**, *13* (14), 3540. https://doi.org/10.3390/cancers13143540.
9. Slamon, D. J.; Leyland-Jones, B.; Shak, S.; Fuchs, H.; Paton, V.; Bajamonde, A.; Fleming, T.; Eiermann, W.; Wolter, J.; Pegram, M.; Baselga, J.; Norton, L. Use of Chemotherapy plus a Monoclonal Antibody against HER2 for Metastatic Breast Cancer That Overexpresses HER2. *N. Engl. J. Med.* **2001**, *344* (11), 783–792. https://doi.org/10.1056/NEJM200103153441101.
10. Cortés, J.; Kim, S.-B.; Chung, W.-P.; Im, S.-A.; Park, Y. H.; Hegg, R.; Kim, M. H.; Tseng, L.-M.; Petry, V.; Chung, C.-F.; Iwata, H.; Hamilton, E.; Curigliano, G.; Xu, B.; Huang, C.-S.; Kim, J. H.; Chiu, J. W. Y.; Pedrini, J. L.; Lee, C.; Liu, Y.; Cathcart, J.; Bako, E.; Verma, S.; Hurvitz, S. A. Trastuzumab Deruxtecan versus Trastuzumab Emtansine for Breast Cancer. *N. Engl. J. Med.* **2022**, *386* (12), 1143–1154. https://doi.org/10.1056/NEJMoa2115022.
11. Horisawa, N.; Adachi, Y.; Takatsuka, D.; Nozawa, K.; Endo, Y.; Ozaki, Y.; Sugino, K.; Kataoka, A.; Kotani, H.; Yoshimura, A.; Hattori, M.; Sawaki, M.; Iwata, H. The Frequency of Low HER2 Expression in Breast Cancer and a Comparison of Prognosis between Patients with HER2-Low and HER2-Negative Breast Cancer by HR Status. *Breast Cancer* **2022**, *29* (2), 234–241. https://doi.org/10.1007/s12282-021-01303-3.
12. Xu, H.; Han, Y.; Wu, Y.; Wang, Y.; Li, Q.; Zhang, P.; Yuan, P.; Luo, Y.; Fan, Y.; Chen, S.; Cai, R.; Li, Q.; Ma, F.; Xu, B.; Wang, J. Clinicopathological Characteristics and Prognosis of HER2-Low Early-Stage Breast Cancer: A Single-Institution Experience. *Front. Oncol.* **2022**, *12*, 906011. https://doi.org/10.3389/fonc.2022.906011.
13. Fu, Z.; Li, S.; Han, S.; Shi, C.; Zhang, Y. Antibody Drug Conjugate: The "Biological Missile" for Targeted Cancer Therapy. *Signal Transduct. Target. Ther.* **2022**, *7* (1), 93. https://doi.org/10.1038/s41392-022-00947-7.
14. Shirman, Y.; Lubovsky, S.; Shai, A. HER2-Low Breast Cancer: Current Landscape and Future Prospects. *Breast Cancer Targets Ther.* **2023**, *Volume 15*, 605–616. https://doi.org/10.2147/bctt.s366122.
15. *Trastuzumab Deruxtecan for Metastatic HER2-Low Breast Cancer - NCI*. https://www.cancer.gov/news-events/cancer-currents-blog/2022/enhertu-her2-low-breast-cancer?utm_source=chatgpt.com (accessed 2025-08-11).
16. Fernandez, A. I.; Liu, M.; Bellizzi, A.; Brock, J.; Fadare, O.; Hanley, K.; Harigopal, M.; Jorns, J. M.; Kuba, M. G.; Ly, A.; Podoll, M.; Rabe, K.; Sanders, M. A.; Singh, K.; Snir, O. L.; Soong, T. R.; Wei, S.; Wen, H.; Wong, S.; Yoon, E.; Pusztai, L.; Reisenbichler, E.; Rimm, D. L. Examination of Low ERBB2 Protein Expression in Breast Cancer Tissue. *JAMA Oncol.* **2022**, *8* (4), 1–4. https://doi.org/10.1001/jamaoncol.2021.7239.
17. Khoury, R.; Saleh, K.; Khalife, N.; Saleh, M.; Chahine, C.; Ibrahim, R.; Lecesne, A. Mechanisms of Resistance to Antibody-Drug Conjugates. *Int. J. Mol. Sci.* **2023**, *24* (11), 9674. https://doi.org/10.3390/ijms24119674.
18. Guidi, L.; Pellizzari, G.; Tarantino, P.; Valenza, C.; Curigliano, G. Resistance to Antibody-Drug Conjugates Targeting HER2 in Breast Cancer: Molecular Landscape and Future Challenges. *Cancers* **2023**, *15* (4), 1130. https://doi.org/10.3390/cancers15041130.

Authors

Cecilia is a high school junior passionate about biology, specifically disease research and bioethics. She is the founder of a community project dedicated to spreading breast cancer awareness and fundraising for local breast cancer foundations. She looks forward to continuing her studies in biomedical sciences at university.

DOI: 10.36838/IJHSR820.22



Marine Biology Research at Bahamas

Unique and exclusive partnership with the Gerace Research Center (GRC) in San Salvador, Bahamas to offer marine biology research opportunities for high school teachers and students.

- Terra has exclusive rights to offer the program to high school teachers and students around world.
- All trips entail extensive snorkeling in Bahamian reefs as well as other scientific and cultural activities.
- Terra will schedule the program with GRC and book the flights from US to the GRC site.
- Fees include travel within the US to Island, lodging, meals, and hotels for transfers, and courses.
- For more information, please visit terraed.org/bahamas.html

Terra is a N.Y. based 501.c.3 non-profit organization
dedicated for improving K-16 education

IJHSR International Journal of High School Research

is a publication of

terra science and education

N.Y. based 501.c.3 non-profit organization
dedicated for improving K-16 education