THE TOWER

SPRING '26



ACKNOWLEDGEMENTS

Faculty and Graduate Reviewers

Jiho Seok
PhD Candidate with the Georgia Tech School of Chemical and Biomolecular Engineering

Dr. Katharine McCann
Director of Undergraduate Research in Neuroscience

Dr. Rheinallt Jones
Assistant Professor in the Department of Pediatric Gastroenterology, Hepatology and Nutrition at Emory University

Dr. M. Talat Odman, Ph.D.
Regents' Researcher Emeritus with the School of Civil & Environmental Engineering at the Georgia Institute of Technology

Faculty Advisors and Mentors

Mac Pitts
Director for Student Media

Dr. Courtney Hoffman
Academic Professional for Undergraduate Research Writing, School of Literature, Media, and Communication

Other Acknowledgements

Brendan Edgecombe,
Melody Lee
Provision of Photos

The Executive Board

Melody Lee
Editor-in-Chief

Kathy Peng
Publications Manager

Ananya Pansuri
Finance Manager

The Editorial Staff

Tvishi Ahluwalia

Zoe Chow

Shreya Iyer

Olivia Lohr

Diya Sharma

Caroline Song

Hanna Zhang

Fall '25 Staff

Mirielle Dogini

Isabel Hollhumer

Letter from the Editor

Dear Reader,

Welcome to the 2025 - 2026 Edition of The Tower, the in-house undergraduate research journal of Georgia Tech. Despite a tumultuous year and a half in the worlds of science, academia, and society at large, I remain grateful to see that undergraduate research at our institution continues to flourish. A brief flip through this year's publication will reveal the colorful range of manuscripts we received in our inbox. Among them are literature reviews, products of on-campus neuroscience classes, and culminations of students' years of collaboration between Georgia Tech and Emory.

Regardless of the nature of the manuscript, these works are reflections of students' desires to learn, to do, and to communicate what they find beautiful. This desire to take action manifests the legacy of Georgia Tech. It turns into a drive for change and a motivator for positive impact, one we are desperately in need of in the coming years. In the words of John Lewis, former civil rights activist and Congressman, for whom our student center is named: "Nothing can stop the power of a committed and determined people to make a difference in our society. Why? Because human beings are the most dynamic link to the divine on this planet."

(From his book, Across That Bridge: Life Lessons and a Vision for Change)

This year, our journal received a record number of submissions. Over the course of the last year, our team of editors has worked tirelessly with the authors to prepare for publication, and I am immensely grateful for their patience and support. I am thankful for our production, content, and business teams for their long hours spent in training and for the many delightful conversations.

I am also grateful to our immediate advisors, Dr. Hoffman and Mac Pitts, who have kindly met with our executive board and respective teams many times throughout these last two semesters. This would not have been possible without your kindness and your guidance, and we can not thank you enough.

Thank you to the authors and their supporting faculty members for their initiative and graciousness in working with the journal. Moreover, we are indebted to the faculty and graduate student advisors. In March, they each received a request to review our manuscripts in their email inboxes, and each and every one of them responded with a resounding "Yes."

And finally, as always, thank you, dear reader, for making it this far and picking up this copy of The Tower. While this journal is only a small example of what committed, determined students are capable of on this campus, we hope you will nonetheless walk away inspired. Look for something that tickles your brain, makes you think, or challenges your perspective. Perhaps you, too, may be driven to try something of your own.

Sincerely,



Melody Lee, Editor-in-Chief

TABLE OF CONTENTS

A Meta-Learning Framework for Automated Selection of PDE Identification Methods (pg. 8)
Pranav Lende, Sung Ha Kang

A Scalable Analog CMOS Architecture for Solving Nonlinear Partial Differential Equations Using Coupled Chaotic Oscillators (pg. 15)
Edwin Zheng, Mrityunjay Krishnakumar, Shrey Jain

Aggression After Traumatic Brain Injury: A Systematic Review of Prevalence, Mechanisms, and Outcomes (pg. 24)
Paige Pieterick

Closing the Coverage Gap: Reforming Georgia's 1332 Waiver to Advance Rural Health Insurance Equity (pg. 32)
Calvin Zhu, Sakthi Priya Ramamoorthy, Sri Dhanya Muppalla, Ria Vittal

Creating New pSB1A3 Plasmids for Fluorescence Research at Georgia Tech Using Generative Techniques for Future Custom Cloning Procedures (pg. 41)
Sanjeev Anand

Developing and Validating Quiet Quitting Intention: A Theory-Based Framework of Intentional Effort Regulation (pg. 50)
Anish Cheekati, Yichen Tang, Christopher Wiese

TABLE OF CONTENTS

Evaluating Air Quality Interventions in Lahore using the Synthetic Control Method (pg. 57)
Sarah Harper, Nicholas Scott, Vianey Guerrero

Evidence and Neurological Mechanisms Behind Sugar Addiction (pg. 66)
Deandra Smith

Freedom to Choose: A Decision-Making Model of Stoic Philosophy (pg. 72)
Ariya Desai

Macrophages and Microbes: Analyzing Inflammation and Repair in the Oral Cavity (pg. 80)
Mirielle Dogini, Keerthi Priya Chinniampalayam Sekar

Menstrual Product Sustainability: Are we there yet? A Look into US and European Union Practices (pg. 88)
Mallika Digavalli, Vibha Iyer, Ophelia Winslett, Marisa Perera, Kayla Carneal, Arina Kucinskaja

Quantifying Polysemanticity in Convolutional Neural Networks (pg. 95)
Nezir Alic, Nayan Ramam, Patrick Ogle, Yuzhou Wang, Kavya Adusumilli, Boyang Gong

The Impact of the 2021 Euro-5 Fuel Standard on PM2.5 Levels in Karachi (pg. 102)
Marika Pescosolido, Dan V. Galea, Jinho Kim

Editor Biographies



A Meta-Learning Framework for Automated Selection of PDE Identification Methods

Pranav Lende

Pranav Lende is a recent graduate of the Georgia Institute of Technology. His interests include scientific machine learning, system identification, and building reproducible tools for noisy spatiotemporal data. This project reflects his focus on designing end-to-end research pipelines that are easy for other researchers to run, verify, and extend.

Abstract

Identifying governing partial differential equations (PDEs) from noisy spatiotemporal data is a central challenge in scientific machine learning. While numerous identification methods exist (including weak-form, sparse regression, and robust approaches), practitioners lack principled guidance on which method to apply for a given dataset. We present a meta learning framework that predicts the best-performing identification method for a data window based on 12 computationally inexpensive features extracted from the raw data, without running any identification algorithm. Using a dataset of 5,786 windows extracted from four canonical PDEs (KdV, Heat, Kuramoto-Sivashinsky, and Transport), we train a Random Forest classifier that achieves 97.06% test accuracy in selecting among four identification methods, driven primarily by sparse-regression dominance on clean simulations. On the held-out test set, 97.06% of predictions achieve the same error as an oracle. Robust methods become essential under high-noise regimes.

1 Problem and Motivation

The task of identifying governing equations from data, often called PDE discovery or system identification, has attracted significant attention in scientific machine learning [Brunton et al., 2016, Rudy et al., 2017]. Given spatiotemporal data $u(x,t)$, the goal is to recover the underlying PDE, typically of the form $u_t = N(u,u_x,u_{xx},...)$, where N is a (possibly nonlinear) differential operator.

A variety of methods have been proposed for this task. Sparse regression approaches like SINDy [Brunton et al., 2016] and PDE-FIND [Rudy et al., 2017] use `1 regularization to identify parsimonious models from a library of candidate terms. Weak-form methods like WeakIDENT [Tang et al., 2023] reformulate the problem using test functions to improve noise robustness. Robust approaches use trimmed regression or `1 loss to handle outliers.

A practical challenge arises when applying these methods: which algorithm should one use for a given dataset? Method performance varies significantly depending on data characteristics such as noise level, spatial resolution, temporal resolution, and the complexity of the underlying dynamics. Running all available methods is computationally prohibitive, particularly for large datasets or in online settings where rapid identification is required.

The challenge of choosing the right algorithm for a given problem instance was first formalized by Rice [1976] as the algorithm selection problem. Recent surveys [Kerschke et al., 2019] show that per-instance algorithm selectors—classifiers trained on cheap instance features to predict the best solver—have achieved strong results in SAT solving, combinatorial optimization, and numerical linear algebra. Our framework applies the same paradigm to PDE identification: inexpensive window-level meta-features route each data instance to the identification method predicted to minimize prediction error, avoiding the cost of running all candidates.

We address this challenge by framing algorithm selection as a supervised learning problem. Given a data window, we extract a small set of features that characterize the data without running any identification method, then use a trained classifier to predict which method will achieve the lowest error. This approach enables efficient, adaptive method selection tailored to each data window.

2 Approach

2.1 Feature Extraction: The Tiny-12 Features

We extract 12 features (referred to as "Tiny-12") from each data window that capture properties relevant to PDE identification without requiring any identification computation. These features are organized into four categories:

1. Derivative statistics (features 0–5): Standard deviations and maxima of spatial and temporal derivatives computed via finite differences: $s(u_x)$, $s(u_{xx})$, $s(u_{xxx})$, $s(u_t)$, $s(u_{tt})$, and $\max|u_t|$.
2. Spectral features (features 6–8): Characteristics of the Fourier spectrum, including energy in low, mid, and high frequency bands, computed via FFT.
3. Signal statistics (features 9–10): Properties of the solution field itself, including s(u) and a nonlinearity ratio that measures the relative magnitude of nonlinear terms.
4. Range feature (feature 11): The dynamic range $\max(u)\min(u)$, which indicates signal amplitude.

These features are intentionally designed to avoid "leakage" from identification outputs. They can be computed in milliseconds, making them suitable for rapid algorithm selection.

2.2 Meta-Learning Selector

We formulate method selection as a multiclass classification problem. For each data window, we compute the Tiny-12 features and use a trained classifier to predict which identification method will achieve the lowest prediction error (measured by the e_2 metric, defined in Section 3.4).

We evaluated six classifiers: Random Forest, Gradient Boosting, K-Nearest Neighbors, Logistic Regression, Support Vector Machine (RBF kernel), and Ridge Classifier. All models were trained using scikit-learn [Pedregosa et al., 2011] with default hyperparameters and a fixed random seed (42) for reproducibility.
The training pipeline consists of: (1) loading the labeled dataset with features and best-method labels; (2) splitting into 80% training and 20% test sets with strat-

ification; (3) standardizing features using z-score normalization (fit on training data only); and (4) training each classifier and evaluating on the held-out test set.

2.3 Classifier Selection Rationale

We evaluated six classifiers spanning four families, each motivated by considerations from the algorithm selection literature [Rice, 1976, Kerschke et al., 2019]:

- **Linear models** (Logistic Regression, Ridge Classifier): Simple baselines with strong inter pretability [Hastie et al., 2009]. If the decision boundary in feature space is approximately linear, these provide efficient and explainable predictions.
- **K-Nearest Neighbors (kNN)** [Cover and Hart, 1967]: This nonparametric method classifies each instance by a majority vote of its k nearest neighbors in feature space, making it a useful reference for detecting locally clustered structure in the meta-feature representation.
- **Support Vector Machine (RBF kernel)** [Cortes and Vapnik, 1995]: A strong general-purpose nonlinear classifier effective in moderate dimensions, capable of capturing complex decision boundaries.
- **Tree ensembles** (Random Forest [Breiman, 2001], Gradient Boosting [Friedman, 2001]): Handle nonlinearities and feature interactions naturally, are robust to feature scaling, and typically achieve strong performance on tabular, heterogeneous feature sets like the Tiny-12 features. A comprehensive comparison of 179 classifiers across 121 datasets found Random Forest to be among the top-performing methods for general classification tasks [Fernadez- Delgado et al., 2014].

2.4 Safety Gate (Future Work)

The current implementation uses a hard classifier that selects a single method. A natural extension is a "safety gate" that runs multiple methods when the classifier's confidence is low. This would involve extracting uncertainty estimates (for example, from Random Forest tree disagreement) and triggering a fallback to top-2 methods when uncertainty exceeds a threshold. We leave the implementation and evaluation of this safety mechanism to future work.

3 Experimental Setup

3.1 PDEs and Data

We generated data from four canonical PDEs, each representing different dynamical regimes:

- **KdV** (Korteweg–de Vries): $u_x + uu_x + u_{xxx} = 0$. This is a dispersive wave equation that admits soliton solutions.
- **Heat**: $u_t = vu_{xx}$. This is a diffusion equation with smooth, decaying solutions.
- **Kuramoto–Sivashinsky (KS)**: $u_t + uu_x + u_{xx} + u_{xxxx} = 0$. This equation exhibits chaotic spatiotemporal dynamics.
- **Transport**: $u_t + cu_x = 0$. This is a linear advection equation with constant wave speed c.

Data for each PDE was simulated on a uniform space-time grid and stored as NumPy arrays. All preprocessing and windowing are fully specified in the repository scripts; the data loading and window extraction steps are deterministic given fixed random seeds.

3.2 Window Extraction

From each simulation, we extracted overlapping windows of size 64×64 (spatial × temporal grid points) with stride 10 in both dimensions. This yielded a total of 5,786 windows distributed as follows:

PDE	Windows	Percentage
KdV	1,734	30.0%
Heat	1,647	28.5%
Transport	1,221	21.1%
KS	1,184	20.5%
Total	5,786	100.0%

3.3 Identification Methods

1. **LASSO**: l_1-regularized regression using scikit-learn's Lasso estimator [Pedregosa et al., 2011] with default regularization.
2. **STLSQ**: Sequentially Thresholded Least Squares, the original SINDy algorithm [Brunton et al., 2016]. Iteratively performs least squares and thresholds small coefficients.
3. **WeakIDENT**: Weak-form identification using test functions and narrow-fit trimming [Tang et al.,

2023].

4. **RobustIDENT**: Robust identification using l_1 loss (least absolute deviation) with l_1 regularization, solved via ADMM.

For each window, we ran all four methods and recorded the resulting error metrics and runtimes.

3.4 Metrics

We used the following metrics to evaluate identification quality and selector performance:

Prediction error (e_2): The normalized residual between the predicted time derivative and the true time derivative:

$$e_2 = \frac{\| \widehat{u}_t - u_t \|_2^2}{\mathrm{Var}(u_t)} \tag{1}$$

where $\widehat{u}_t$ is the prediction from the identified model and u_t is the numerically computed time derivative.

Oracle error: For each window, the oracle error is $\min_m e_2^{(\mathrm{m})}$, the minimum error across all methods m.

Regret: The difference between the selector's chosen method error and the oracle error:

$$\mathrm{Regret} = e_2^{(\mathrm{selector})} - e_2^{(\mathrm{oracle})} \tag{2}$$

Zero regret indicates the selector achieves the same error as the oracle (i.e., the minimum possible e2 for that window).

Zero-regret rate: The fraction of windows where regret equals zero.

3.5 Reproducibility

Code and data repository: https://github.com/ALLEN-DE123X/ident-meta

All experiments used fixed random seeds (42) for train-test splitting and model initialization. Window extraction follows a deterministic grid stride. The full pipeline can be reproduced using the following commands:

```
python scripts/run_all_methods.py # ~25 min
python scripts/train_models.py # ~2 min
python scripts/generate_figures.py # ~1 min
```

4 Results

4.1 Model Comparison

Table 1 shows the performance of six classifiers on the method selection task. Random Forest achieved the highest test accuracy at 97.06%, followed by Gradient Boosting at 95.68%. We report 5-fold cross-validation accuracy using scikit-learn's default classification cross-validation setting [Pedregosa et al., 2011]. The gap between the cross-validation mean and the single held-out test split suggests substantial variance across

Table 1: Classifier comparison for PDE method selection.

Model	Test Accuracy	5-Fold CV Mean ± Std
Random Forest	**0.9706**	0.879 ± 0.125
Gradient Boosting	0.9568	0.876 ± 0.122
KNN (k=5)	0.9499	0.872 ± 0.116
Logistic Regression	0.8946	0.884 ± 0.109
SVM (RBF)	0.8869	0.863 ± 0.127
Ridge Classifier	0.8800	0.865 ± 0.107

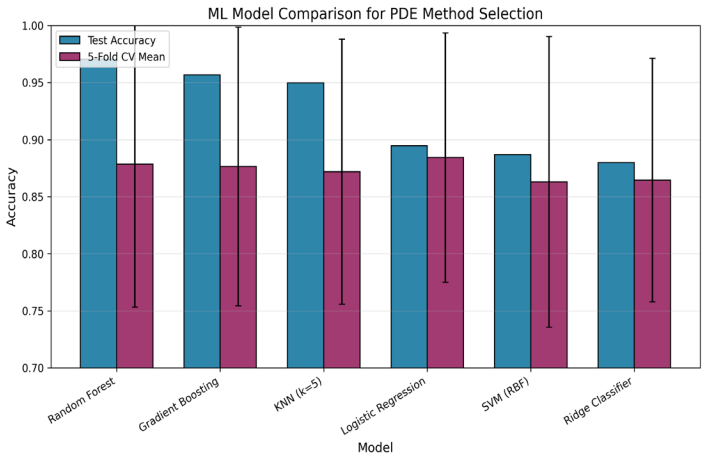

Figure 1: Comparison of classifier accuracy for method selection. Random Forest achieves the highest test accuracy, though several methods perform competitively.

4.2 Regret Analysis

Using the Random Forest selector, we computed regret on both the held-out test set and the full dataset:

Held-out test set (1,158 windows):

- Zero-regret count: 1,124 of 1,158 windows (97.06%)
- Mean regret: 0.0010
- Max regret: 0.4396

Full dataset, in-sample (5,786 windows):

- Zero-regret count: 5,752 of 5,786 windows (99.41%)
- Mean regret: 0.0002

The higher in-sample zero-regret rate is expected since this includes training data. The held-out test results provide a more conservative and realistic estimate of performance on unseen windows. Note that zero regret means the selector's chosen method achieves the same e_2 as the oracle; it does not require the predicted label to match the oracle label, since ties in e_2 across methods are possible.

Table 2: Regret summary for the Random Forest selector. Regret is defined as $e_2^{(\text{selector})} - \min_m e_2^{(\text{selector})}$; zero regret indicates oracle-level error.

Evaluation	Zero-Regret %	Zero-Regret Count	Mean Regret	Max Regret
Test (1,158 windows)	97.06%	1,124	0.0010	0.4396
In-sample (5,786 windows)	99.41%	5,752	0.0002	0.4396

Figure 2 shows the cumulative distribution of regret values. The steep rise at zero indicates that the vast majority of predictions exactly achieve oracle-level error.

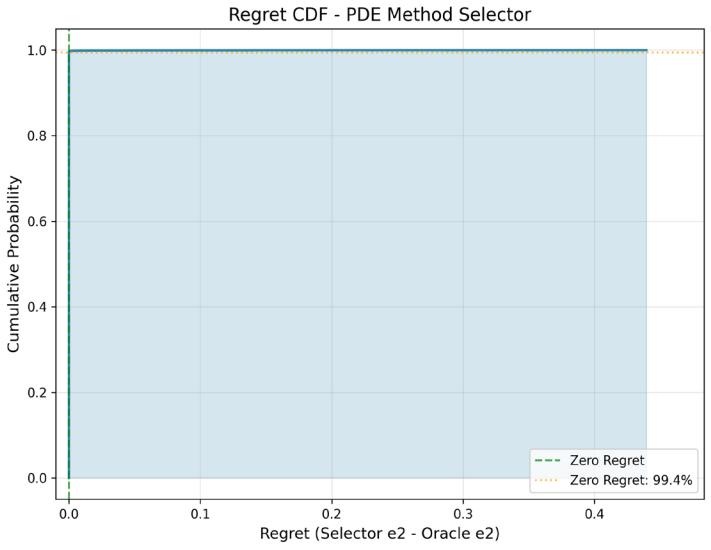

Figure 2: Regret CDF evaluated in-sample over the full dataset. The curve rises sharply at zero regret because 99.41% of windows achieve oracle-level error.

4.3 Confusion Matrix

Figure 3 shows the confusion matrix for the Random Forest classifier on the test set. Due to extreme class scarcity in the full dataset (only 4 WeakIDENT and 2 RobustIDENT samples total), the held-out test split contains 729 LASSO, 428 STLSQ, 1 WeakIDENT, and 0 RobustIDENT samples. Consequently, the test-set evaluation largely reflects a binary distinction between LASSO and STLSQ. The model achieves high precision and recall for both dominant classes.

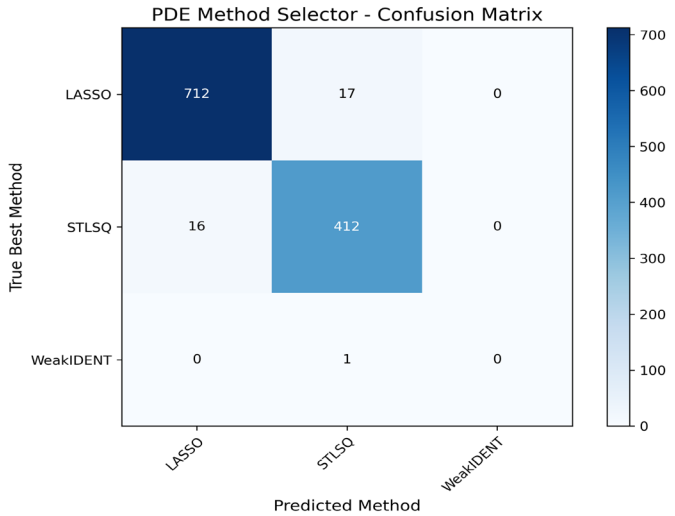

Figure 3: Confusion matrix for method selection on the test set. Classes with zero samples in the test set (RobustIDENT) are not shown. Strong diagonal dominance indicates accurate classification for the two dominant classes.

4.4 Feature Importance

Figure 4 shows the Random Forest feature importances. The spatial derivative features $s(u_{xx})$ (feature 1), $s(u_{xxx})$ (feature 2), and $s(u_x)$ (feature 0) are most predictive, in that order. These derivative statistics capture the spatial complexity of the dynamics, which strongly influences method performance.

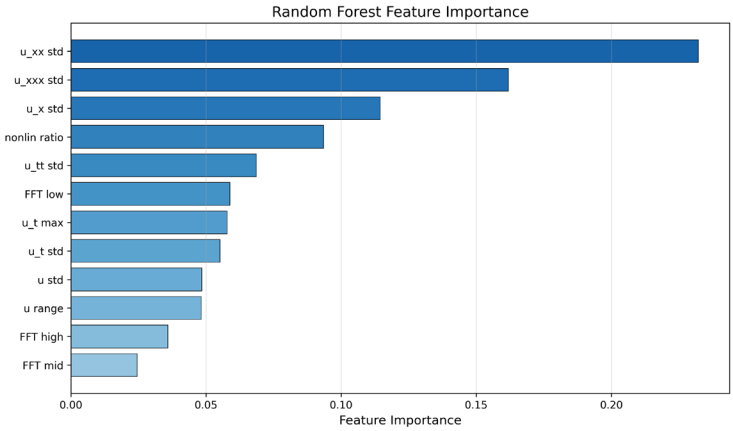

Figure 4: Feature importance for the Random Forest

selector. Spatial derivative statistics are most predictive of optimal method choice.

4.5 Best-Method Distribution

Figure 5 shows the distribution of best methods across the dataset.

The distribution is notably skewed with exact counts: LASSO is optimal for 3,644 windows (63.0%), STLSQ for 2,136 windows (36.9%), WeakIDENT for 4 windows (0.07%), and RobustI DENT for 2 windows (0.03%). This extreme imbalance reflects the favorable conditions of our clean synthetic data, where fast sparse regression methods perform well. The train-test split uses stratification (stratify=y), but with only 6 samples in the minority classes, these classes are inherently difficult to learn; the classifier effectively operates as a binary LASSO-vs-STLSQ selector in this regime. A naive baseline of "always select LASSO" would achieve 63% accuracy, compared to the selector's 97.06%, a gain of 34 percentage points. The selector learns to correctly identify the 37% of cases where STLSQ outperforms LASSO.

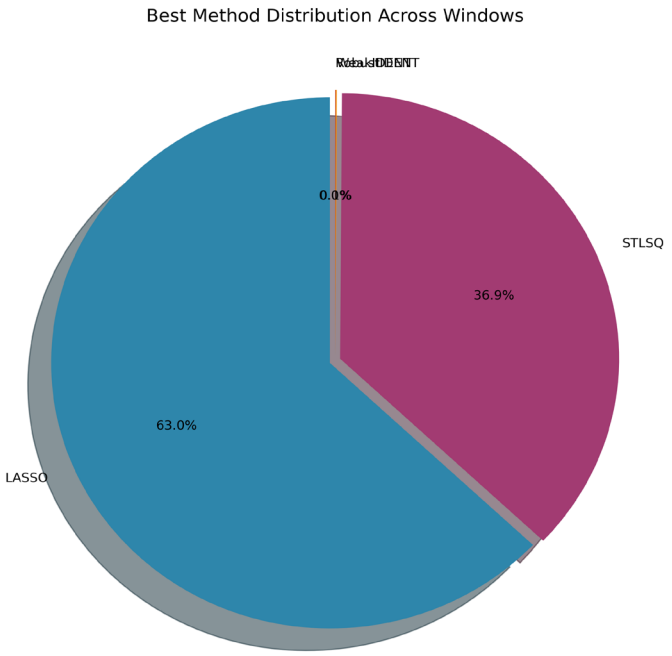

Figure 5: Distribution of best-performing methods across windows. LASSO and STLSQ dominate on this clean synthetic dataset.

4.6 Noise Stress Test

To investigate how method performance changes under data corruption, we conducted a noise stress test on a subset of 200 randomly sampled windows. We applied controlled corruption to each window before running the identification methods: additive Gaussian noise (2% and 5% of signal standard deviation) and outlier corruption (1% and 3% of values replaced with 5–10× spikes).

Figure 6 shows how the best-method distribution shifts under each corruption setting. Key findings from this experiment (200-window subset):

- On the clean subset, LASSO wins 56.5% and STLSQ wins 43.5% of windows (similar ordering to the main dataset, though percentages differ due to random sampling).
- Under 2% Gaussian noise, RobustIDENT becomes optimal for 20.5% of windows, while LASSO drops to 44.5% and STLSQ to 35.0%.
- Under 5% Gaussian noise, LASSO captures 53.0%, STLSQ 28.0%, and RobustIDENT 19.0%.
- Under 1% outlier corruption, LASSO captures 69.0%, STLSQ 12.0%, and RobustIDENT 19.0%.
- Under heavier corruption (3% outliers or combined 5% noise + 1% outliers), LASSO domi nance returns (83–87%), suggesting that all methods degrade similarly and LASSO remains a reasonable default.

These results confirm that robust methods become competitive under moderate adverse condi tions, motivating the need for adaptive method selection rather than a one-size-fits-all approach.

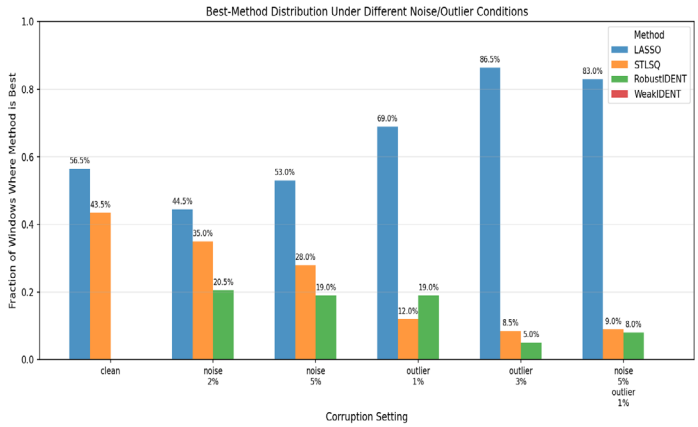

Figure 6: Best-method distribution under different corruption settings. Under 2% and 5% Gaussian noise, mpRobustIDENT becomes the optimal method for approximately 20% of windows, demonstrating that ro-

bust methods gain importance under noisy conditions.

5 Limitations and Future Work

Skewed class distribution: On clean synthetic data, LASSO and STLSQ dominate, accounting for 99.9% of best-method selections (5,780 of 5,786 windows). WeakIDENT and RobustIDENT collectively appear as optimal in only 6 windows. This reflects the favorable conditions of our synthetic experiments (low noise, sufficient resolution). Our noise stress test demonstrates that RobustIDENT captures up to 20% of windows under 2% Gaussian noise, confirming that robust methods become important under adverse conditions. Future work should systematically evaluate method performance across a noise-resolution parameter sweep.

Stratification brittleness: The train-test split uses stratification, but with only 6 total samples in the minority classes (WeakIDENT: 4, RobustIDENT: 2), the classifier cannot reliably learn to select these methods. Performance on minority classes should be interpreted cautiously.

PySINDy and WSINDy: These PySINDy-based methods are available in the Docker environ ment but not in the standard installation due to dependency conflicts. Extending the comparison to include these methods is straightforward using the Docker workflow.

Cross-PDE generalization: Our current evaluation uses random train-test splits. A leave-one PDE-out cross-validation would provide stronger evidence of generalization to unseen equation types.

Safety gate: Implementing uncertainty-based fallback to top-2 methods would improve robust ness when the classifier is unsure.

Runtime-aware selection: The current objective minimizes prediction error. Incorporating runtime into the objective (for example, error plus a penalty on compute time) would enable tradeoffs between accuracy and speed.

Conclusion

We presented a meta-learning framework for automated PDE identification method selection. By extracting 12 inexpensive features from raw data, a Random Forest classifier predicts with about 97% test accuracy which of four identification methods will perform best. On our held-out test set of 1,158 windows, 97.06% of predictions achieved the same error as an oracle that always selects the best method. This framework enables efficient algorithm selection for PDE discovery, avoiding the computational cost of running all methods on every data window. Future work will extend the approach to noisy data regimes, implement uncertainty-based safety mechanisms, and evaluate cross-PDE generalization.

Acknowledgements

The author thanks Dr. Sung Ha Kang for proposing the research direction and providing mentorship and feedback throughout the semester. All implementation, experiments, and writing in this manuscript were completed by the author. The author also thanks the Georgia Institute of Technology for providing an environment and resources that supported this work.

Access references for this paper using the QR code below.



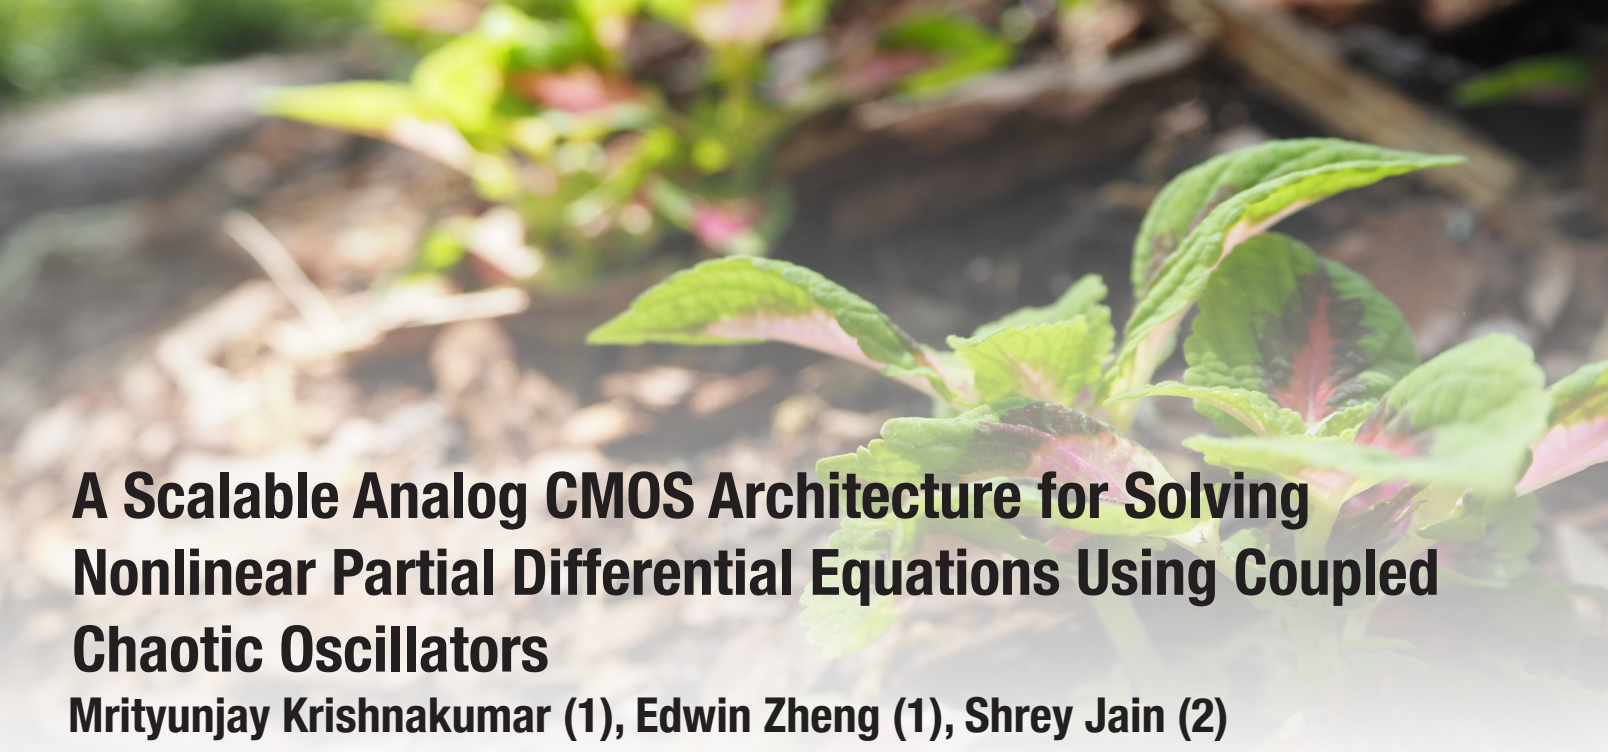

A Scalable Analog CMOS Architecture for Solving Nonlinear Partial Differential Equations Using Coupled Chaotic Oscillators

Mrityunjay Krishnakumar (1), Edwin Zheng (1), Shrey Jain (2)

Mrityunjay Krishnakumar is a first year electrical engineering major. Audio engineering and analog signal processing are two areas that he is extremely passionate about. His area of interests focus on where music technology and hardware engineering meet and how we can create new tools for generating and producing sound. As a professional, his goal is to be an Audio Hardware Engineer, developing the next generation of synthesizers, effects units, and other professional grade audio equipment.

Edwin Zheng is pursuing a Bachelor of Science in Electrical Engineering at the Georgia Institute of Technology and will graduate in 2028. His focus lies in analog ASIC design and the use of analog technologies to advance signal processing and computing as well as progressing consumer technologies, using theoreticals to create practical electronic designs. He is also interested in the fusion of digital and analog elements to create devices that help bridge together the two domains thus helping them achieve the highest usability in the ever-moving status quo of modern computational technologies. Finally, he hopes to use analog circuits as accelerators for computations currently done in the digital domains in order to create faster, more efficient chips that will benefit the general consumers.

Shrey Jain is a first year computer science major at Georgia Tech. He is interested in nonlinear dynamics, analog computing, and computational modeling. Specifically, he is interested in how physical hardware can function as energy-efficient accelerators for scientific computation. He is fascinated by architectures that bridge computer science and electrical engineering by utilizing dynamical systems as computational substrates. His goal is to investigate how continuous-time analog systems can provide for newer approaches to solving certain nonlinear PDEs.

Abstract

This paper presents an analog computing approach that can be scaled to solve nonlinear partial differential equations (PDEs), utilizing arrays of chaotic oscillators based around standard CMOS. Chaotic Oscillators are analog circuits that can be modeled by nonlinear PDEs and operate continuously with voltage and current. When these oscillators are coupled, the resulting collective interaction between them will mimic the behavior of a complex system. This allows the array to perform the solution of PDEs by leveraging the inherent physical circuit laws to achieve massive parallelism while performing computations at significantly less than the energy required by digital solvers. Since the architecture processes information through continuous time-physical laws instead of the discrete step-by-step manner of digital processors, it avoids the von Neumann Bottleneck that limits traditional digital processors. The proposed system uses a 2D grid of chaotic unit cells built from modified versions of Chua's Circuits that are configured to produce chaotic behavior analogous to the Lorenz Butterfly Effect. The chaotic unit cells are coupled together via programmable transconductance elements to simulate reaction diffusion systems and wave propagation phenomena.

This work presents a theoretical framework describing how such a CMOS-integrated chaotic array would be designed, coupled, and parameterized to approximate solutions to target PDEs while maintaining robustness to process variation. By analyzing the expected collective dynamics of the coupled oscillators under this framework, the proposed architecture is shown in principle to provide a viable high-performance option for solving high-dimensional computational problems.

1. Introduction

In the last couple of decades, the development of theoretical designs and the circuit implementation of chaotic systems as well as the generation of multi-scroll attractors has become an increasingly popular subject due to its ability to be used in a variety of different ways in secure communication, signal processing, and cryptography [14]. A large improvement in the study of chaotic systems is the creation of the multi-scroll attractor. The multi-scroll attractor is a type of chaotic attractor that can be described by a complex phase space topology with a number of localized regions of chaotic wandering that resemble spirals. This increase in the complexity of the chaotic attractor increases the pseudo-randomness and structural diversity of the chaotic system. Unfortunately, one of the main barriers to integrating chaotic systems with standard complementary metal-oxide-semiconductor (CMOS) technology to produce multi-scroll attractors with electronic circuits built using discrete components is the presence of parasitic capacitance within the CMOS technology, sub-threshold leakage current, and the amount of silicon area that is needed to implement each of the passive components that are typically present in chaotic systems [9], all of which can severely distort the precise piece-wise linear function that is needed to support multi-scroll dynamics.

To appreciate why it is necessary to develop these new chaotic systems, you must first examine how we got to where we are today with computing. In 1965 Gordon Moore (co-founder of Intel) predicted that there would be an exponential increase in the number of transistors on an Integrated Circuit (IC), that this doubling would occur at least once per year over the next decade with little or no increase in cost; five years later he doubled the rate of this doubling to every two years [15]. As a result of this philosophy, companies have been working to put ever smaller transistors on ever smaller chips at ever lower costs. Today, however, digital transistor sizes are almost at their minimum limit — a well-known problem called the end of Moore's law [2]. Digital transistor sizes are now very close to atomic dimensions and geometric shrinkage is therefore no longer possible without causing significant problems due to quantum tunneling and signal degradation. Additionally, as the distance between logic gates decreases so does the heat generated when a chip is used. Supercomputer-sized facilities are extremely high density with thermal energy, so they are not only operated but are also required to equip cooling towers that are so huge that they evaporate millions of gallons of freshwater air per day and incur exorbitantly high economic costs, aside from being the cause of depletion of environmental resources [16]. Scaling down the high component count is one way to deal with high heat issues.

Chaos can be seen as an area that has a major overlap between system design, physical interfacing, and integrated circuit engineering; hence, there is also a considerable computational cost related to chaotic nonlinear system simulation on conventional computers [7]. Mathematical operations in digital systems are performed by finite floating-point binary approximations. Because chaotic systems are exponentially sensitive to the initial conditions—the butterfly effect—quite the small initial rounding errors of digital computers work as the continuous perturbations. These truncation errors build up quickly, and the simulated trajectory of the system diverges completely from the actual mathematical solution unless very small, computationally expensive integration steps are applied [17].

The move to "beyond Moore" architectures has prompted the development of new non-Boolean computing models that do not use digital logic elements but rather can do computations through the nonlinear interactions of oscillating components [2], [3]. The 'Beyond Moore' model can be considered an ideological alternative to the law of technological geometric scaling; it encompasses the new physical characteristics as well as the parallel networks which determine the new behavior of the computational systems. In contrast to a classical logic element (AND gate or OR gate) that simply evaluates a binary truth table through either the switching of electrical charge on or off, non-Boolean computation models (such as oscillators and reservoir computers) work based on holistic processing of

information rather than based upon direct assignment of input data to table entries. They don't derive their solutions to mathematical problems one algorithm at a time- instead the network of components breathes energy into the air and by the process of physical exchange of energy it completely settles into a state where the energy in the whole network is minimal and they all are in a synchronized way.

The reason Chaotic Chua Circuits are so interesting is because they are very significant. In 1983, Leon O. Chua developed the Chua Circuit which is an example of an extremely simple autonomous electric network which exhibits chaotic behavior [18]. It achieves this through Chua's diode, a locally active nonlinear resistor, along with three energy-storage elements that generate those surprising voltage oscillations. By linking multiple Chua circuits with resistors or amplifiers, engineers have successfully created coupled chaotic oscillators. The core focus of this study revolves around the fundamental phenomena of these coupled chaotic oscillators, particularly exploring how they interact to achieve synchronization and their potential for processing spatial information [13].

This work introduces a novel model of analog computing where chaotic coupled oscillators act as a single computational unit. Previously, analog computing faced challenges due to the need for large inductors to generate chaotic oscillations. However, researchers have demonstrated that it's possible to create analog versions of Chua's circuit without any inductors by utilizing a current feedback operational amplifier (CFOA) [5], [19]. CFOAs are a great resource due to their versatility compared to the standard components (voltage-domain) that most engineers work with; this device acts as an "earthed" capacitor and resistor in place of grounded capacitors and resistors when creating inductors on a chip. These studies have demonstrated that CFOA implementations will function equivalently as those using external coils at high frequencies thus decreasing the space needed on the chip.

Recent studies have developed a smaller number of physical devices to create chaotic and hyperchaotic circuits with fewer than 10 total components, and coupled some of these circuits together using current-mode operation as well [1], [11]. The use of current-mode operation also simplifies the process of connecting circuits in large-scale systems. A dual output second-generation current conveyor (DO-CCII) is a type of active analog circuit that produces a mirrored current at its two outputs in response to a single voltage input [20]. Because a DO-CCII inherently mirrors current, it eliminates the need for expensive, space-consuming external resistors, and allows for much simpler routing of signals through the dense network of oscillators.
No matter the fact that the digital versions of analog chaotic systems have come a long way, they still rely on traditional analog simulators which are still suffering from high power consumption. This paper is aimed at the introduction of dynamic threshold metal oxide semiconductor (DTMOS) technology as a theoretical way to cut down the chaotic core size and the power consumption that comes with it. Scientific literature recommendations enforce that the massively parallel analog grids should follow the stringent thermal design limit; it takes ultra-low supply voltages to prevent thousand coupled nodes from thermal runaway during operation. The DTMOS circuit delivers this efficiency by means of a design that ties the gate terminal of the transistor directly to its body [21]. This way, the internal threshold voltage is dynamically lowered while the device remains active, that is why this configuration achieves high-speed operation at low supply voltages - on the other hand, the threshold voltage is raised when the device is not active to inhibit leakage current. The idea is to present a non-inductor circuitry which is beyond DTMOS-based DXCCII architectures for low power operations targeting the microwatt (1.58 μW) and not the milliwatt range for power consumption. The CMOS chaotic unit cells can also be physically arranged in a 2D grid. The coupling within the grid represents a physical medium that evolves towards a solution to a partial differential equation (PDE) due to its continuous nature [10]. Due to the uniformity and continuity of the nodes within the grid and their connection to their neighbors, it is acceptable to make the continuous mathematical assumption of uniform coupling. The uniform nearest neighbor connections between nodes create a spatial Laplacian operator that is analogous to reaction-diffusion systems based on the difference in voltage between each pair of neighboring nodes [12].

Within this theoretical framework, the robustness to PVT variations and the synchronization necessary for coupled computation represent intrinsic characteristics

of the collective oscillator dynamics of the units versus characteristics that would need to be verified by the use of physical fabrication or circuit level simulation. PVT values represent the operation conditions of a manufactured device, as well as the tolerance limits of the manufacturing process of that device [22].

Table 1

Parameter	Description	Impact on CMOS Circuits
Process (P)	Microscopic variations in manufacturing (e.g., doping, lithography)	Alters threshold voltages and transistor channel lengths
Voltage (V)	Spatial and temporal power supply fluctuations	Causes IR drops (voltage drops) and affects gate switching speeds
Temperature (T)	Thermal gradients resulting from active power dissipation	Modifies carrier mobility and increases sub-threshold leakage

The ambient pressure has been intentionally left out of this standard variation analysis, since all integrated circuits are housed within their protective encapsulations. Therefore, normal atmospheric changes cannot cause any mechanical stress upon the die, nor can they affect its electrical bandgap properties. Although there may be some unavoidable microscopic differences in PVT mismatch between the oscillators on the die, when the transconductance coupling between the oscillators exceeds a certain critical value, the entire network of oscillators will spontaneously lock into phase. In other words, the required synchronization for computing over space will emerge through the macroscopic behavior of the system. Since the system's micro-scale PVT flaws of the individual components will average out due to the natural macroscopic behavior of the system, the spatial computations will be reliable.

2. Prior Work and Limitations

2.1 Classical Analog PDE Solvers

In order to approximate Laplacian operators, prior analog PDE machines made use of resistor meshes and op-amp networks. Even though these systems succeeded in solving the Poisson and heat diffusion equations, they were limited in terms of linear PDEs since they were unable to implement arbitrary nonlinear reaction functions [15].

The Poisson equation comes up in electrostatics based on Gauss's law under steady-state charge distributions, whereas the heat diffusion equation arises from Fourier's law of thermal conduction and conservation of energy. Their solutions can be created from linear combinations of basis responses and so these equations are linear and governed by superposition. They are well suited for solving such linear elliptic and parabolic PDEs in steady state since resistor networks implement discrete Laplacians through Kirchhoff's Current Law [15]. At times, nonlinearities were approximated with saturating operational amplifiers or semiconductor diodes whose nonlinear current–voltage characteristics led into limited thresholding behavior. These elements gave fixed nonlinear results that didn't have precise control and lacked programmability, preventing accurate implementation of general nonlinear reaction terms. This limitation is crucial since many relevant PDEs—including reaction–diffusion systems, chaotic field equations, and excitable media—need spatially distributed nonlinear interactions that can't be captured by linear resistor networks.

2.2 Chaotic Oscillators in CMOS

Chua's circuit has been analyzed and implemented with MOSFETs, piecewise-linear resistors, as well as inductor emulators. Recent work has allowed for inductor-less realizations using microwatt-level power consumption, allowing for integration of large arrays of oscillators. By replacing bulky spiral inductors with active synthetic inductors, this scalability has been enabled, lowering the supply voltages with subthreshold and body-biasing techniques, and utilizing current-mode architectures that ensure there's low passive component count [1,6,7]. However, current research places emphasis on random number generation, chaos theory, and secure communication. No current CMOS chaotic oscillator arrays have been configured as continuous-time nonlinear PDE solvers yet. This is significant because chaotic oscillators obey nonlinear differential equations, indicating that large coupled arrays could be physical computational substrates for solving spatially distributed dynamical systems, along

with being signal generators.

2.3 Hardware PDE Solvers

Even though we can see that many neuromorphic or accelerators inspired by physics exist, none of them portray a CMOS-integrated chaotic lattice as a PDE solver [10,16]. Integrating such functionality into CMOS would allow for on-chip continuous-time computation of nonlinear PDEs. This could potentially eliminate algorithmic iteration present in digital solvers and also reduce energy consumption. This gap is notable since chaotic CMOS circuit design and analog computing research have progressed independently, despite their shared foundation in terms of nonlinear dynamical systems theory. This research work helps bridge this gap by proposing a lattice of coupled chaotic unit cells whose collective dynamics implement discretized nonlinear PDEs. Moreover, this unifies chaotic circuit theory with analog computational modeling.

3. Reaction-Diffusion Theory and Circuit Mapping

3.1 Reaction Diffusion PDE

A scalar reaction-diffusion PDE is defined as:

$$\left(\partial u/\partial t\right) = D\nabla^2 u + f\left(u\right)$$

where:

- the diffusion term $D\nabla^2 u$ smooths the local variations
- the nonlinear term $f(u)$ introduces pattern formation, bistability, or oscillation

The classic cubic reaction law:

$$f\left(u\right) = au - bu^3$$

produces Turing instabilities as well as spatial patterning.

In this equation, u is the state variable, a governs linear growth, and b limits amplitude with nonlinear saturation.

3.2 Finite- Difference Discretization

Using second-order finite differences, the PDE turns into a system of coupled ODEs:

$$\left(du_i/dt\right) = D(u_{i+1} - 2u_i + u_{i-1}) + f(u_i)$$

where each u_i represents the field (discrete spatial point). The grid coupling term $(u_{i+1} - 2u_i + u_{i-1})$ acts as a discrete Laplacian. This is known as a discrete Laplacian since by using the nearest-neighbor differences derived from finite-difference methods, it approximates the second spatial derivative [15].

3.3 CMOS Mapping

Let's say that Vi denotes the node voltage that represents ui. In that case, with a state capacitor C, the node dynamics become:

$$C(dV_i/dt) = g(V_i) + g_m(V_{i+1} - 2V_i + V_{i-1})$$

where:

- $g(V_i)$is the nonlinear Chua-diode current that implements the reaction term
- g_m is the effective transconductance of the coupling path
- $D = g_m/C$

This sets up a one-to-one mapping between the circuit and the PDE. This mapping follows from Kirchhoff's Current Law, where the sum of local nonlinear current and neighbor coupling current equals the capacitor voltage derivative, which matches up the structure of the discretized PDE. The CMOS array acts as a continuous-time physical solver for the discretized nonlinear PDE.

4. Methodology: Scalable CMOS Implementation of Coupled Chua-Circuit Arrays for Massive Parallel Solution of Non-Linear PDEs

4.1 Overview and Design Philosophy

The design of this proposed system considers a discretized form of a reaction-diffusion partial differential equation (PDE) as a collection of equivalent chaotic unit cells. These unit cells have similar state variable characteristics of the voltages of the first capacitor $C1$ of each Chua circuit. By linearly coupling neighboring unit cells through resistive (conductive) interconnections, the diffusion term is created, which enforces current proportional to voltage differences between adjacent nodes. This produces a current term proportional to (Vj – Vi). By Ohm's law and Kirchhoff's Current Law, implementing the discrete Laplacian

operator described in Section 3 [15]. The nonlinear reaction term is created by the chaotic behavior of the Chua diode within each unit. Each of the unit cells evolves over continuous time, so the overall solution for the spatial domain occurs in one analog time-step; thus offering speed independent of grid size in analog operation. The implementation of this method utilizes three previously developed CMOS building blocks: (i) a Chua oscillator that does not require an inductor, implemented with a synthetic inductor, (ii) a compact, piece-wise-linear (PWL), Chua diode implemented with a current-mode active block, and (iii) a programmable MOS-triode coupling network. These building blocks have been demonstrated in CMOS implementations of chaotic oscillators and inductor-less Chua circuits [5,6,8,9,11].

4.2 Unit-Cell Design

4.2.1. Synthetic Inductor (Active Inductor)

A conventional Chua circuit requires a physical inductor that is prohibitive for large-scale integration. The synthetic inductor was implemented with a gyrator-C topology with a dual-X second-generation current conveyor (DX-CCII) and a grounded capacitor. A gyrator is an active two-port network that basically converts a capacitor into an equivalent inductance with impedance inversion, getting rid of the need for a large spiral inductor on-chip C_L The DX-CCII provides high bandwidth, low-noise current-mode operation, and its port relationships enable the impedance transformation $Z_{in} = jwL_{eq}$ where $L_{eq} = R_a R_b C_L$ with R_a and R_b being the two resistors attached to the DX-CCII ports. The synthetic inductor occupies only a few micrometres of silicon area, eliminating the need for a spiral inductor and permitting thousands of cells per chip. On-chip spiral inductors in standard CMOS processes usually occupy hundreds of square micrometers and portray low quality factors because of substrate losses, making them area and power inefficient for dense integration. Substituting them with active synthetic inductors reduces the area by a lot and also leads to improved scalability in terms of big oscillator arrays [5,9]. This is crucial since spiral inductors take up a significant silicon area, while synthetic inductors only use active capacitors and devices, improving integration density. Experimental validation of this approach has shown a functional inductor-less Chua oscillator operating at a bias voltage of 0.2 V and consuming 1.58 μW.

4.2.2. CMOS Chua Diode (Non-Linear Resistor)

The chaotic behavior is generated by a PWL negative-resistance characteristic. Using a single DX-CCII together with three precision resistors, the inner negative-slope m1 and outer positive-slope m0 can be programmed by selecting the transistor aspect ratios in the DTMOS implementation. The outputted I-V curve reproduces the well-known “N-shaped” Chua diode and is capable of supporting a double-scroll attractor for a wide range of bias voltages. A second design using a voltage-differencing transconductance amplifier (VDTA) provides an operating frequency in excess of tens of MHz and consumes just 0.243 mW. Both designs are fully compatible with a 0.18 micron CMOS process allowing for sub-micron device sizes and ultra-low power operation.

4.2.3. Low-Voltage, Low-Power Operation

Current mode active blocks utilizing dynamic-threshold MOS (DTMOS) transistors were utilized to implement both the DX-CCII and VDTA stages in order to support operation at supply voltages as low as 0.2 V. The utilization of the DTMOS body-effect provides a reduced threshold voltage when the gate is high, providing increased drive capability in the sub-threshold region, while maintaining linearity in the analog domain. This approach limits the static power dissipation to the uW level and allows the entire unit-cell to be powered from a single low-voltage rail, which is required for densely packed arrays.

4.3 Coupling Network: Mapping the Diffusion Term

4.3.1. Resistive Coupling Using MOSFETs in Triode

The Laplacian discretization

$$D\nabla^2 V \approx \frac{D}{\Delta x^2} \sum_{nbr} (V_j - V_i)$$

was implemented by connecting the nodes of adjacent cells together with MOSFETs biased in the deep-triode region. The channel resistance of a MOSFET operating in the deep-triode region is:

$$R_{triode} \approx \frac{1}{\mu C_{ox}, \frac{W}{L}, (V_{GS} - V_{TH})_b}$$

and can be adjusted dynamically by controlling the gate voltage VGS , which results in a programmable diffusion coefficient.

$$D \propto \frac{1}{R_{triode}}$$

This coupling scheme represents a current-mode coupling scheme similar to that presented for Chua arrays, where the availability of the inductor current supports the creation of a current-based diffusion term without the need for additional voltage buffers. [5].

4.3.2. Programmable Conductance Matrix

A digital-to-analog converter (DAC) per coupling transistor allows the diffusion coefficient to be defined at each point along a row/column (i.e. heterogeneous diffusion profile) which can be used to define either spatially varying diffusion coefficients or anisotropically diffusing particles. The DACs may be shared among multiple rows/columns to minimize area, and the control registers are addressable via an on-chip RISC-V processor allowing for real-time configuration of the diffusion coefficients.

4.4 Array Architecture and Boundary Conditions

4.4.1. 2-D Cartesian Grid

Cells are tiled in a regular 2-D grid with a pitch of ≈ 10 µm, determined by the DX-CCII footprint and coupling transistor dimensions. The array can scale to N x N where N is limited only by the die area (e.g., a 5 mm × 5 mm die can host > 10^6 cells). Each cell's state is sampled by a column-parallel ADC for read-out or for feedback to higher-level digital control loops.

4.4.2. Boundary Implementations

Neumann (Zero-Flux) Boundaries – Edge cells leave the coupling MOSFETs open-circuit, yielding zero current flow and thus implementing $\partial V/\partial n = 0$.

Dirichlet (Fixed-Value) Boundaries – Edge cells are tied to a reference voltage through a low-impedance MOSFET switch, fixing the node voltage to the desired Dirichlet value.

These configurations are identical to those used in coupled-oscillator analog computers, where the termination impedance determines the effective boundary condition [3] .

4.5 Parameter Tuning and Reaction-Term Control

The reaction term $f(V)$ is defined by the shape of the Chua diode's PWL characteristic. By adjusting the inner and outer slopes$(m_1,\ m_0)$ through transistor sizing, the non-linear function can emulate a wide class of reaction kinetics (e.g., cubic, logistic, or FitzHugh-Nagumo type). This is possible since the programmable slopes and breakpoint voltages of the piecewise-linear Chua diode allow for the estimation of polynomial nonlinearities that characterize many reaction terms. In addition, a global bias current injected into the DX-CCII can shift the operating point, allowing fine-tuning of the chaotic regime to match the desired non-linear function.

4.6 Simulation and Validation Flow

To acquire parasitic resistances and capacitances of the DX-CCII, VDTA, and coupling MOSFETs, device-level extraction was performed with a 0.18 µm CMOS PDK. After that, a SPICE macro-model of a single unit cell was built including the synthetic inductor, Chua diode, and coupling transistors. In order to verify double-scroll attractors and diffusion coupling, transient simulations were utilized. The SPICE model was converted to a Verilog-A behavioral model in order to enable Monte Carlo analysis of process and temperature variation [8]. After that, an 8×8 array was co-simulated in Cadence Virtuoso AMS and compared against a finite-difference numerical solution to evaluate convergence speed and accuracy.

Table 2

Metric	Target (Theoretical)	Rationale
Operating Frequency	$10 - 30$ MHz (VDTA) or 1-5 MHz (DX-CCII)	Sufficient enough to resolve fast reaction dynamics while also ensuring that power is low [4, 5]

Power per Cell	$\leq 2\mu W$ (DTMOS)	Showed in the low-voltage inductor-less Chua implementation [1, 2]
Area per Cell	$\leq 30\mu m^2$	Synthetic inductor and DX-CCII occupy $< 15\mu m^2$; coupling MOSFET adds $< 5\mu m^2$
Diffusion Coefficient Range	$10^{-3} - 10^2 mm^2/s$ (tunable via gate bias)	Direct control of triode resistance gives > 3 decades of tunability
Solution Latency	$< 1mm^2$ for a 256 x 256 grid (set by circuit settling)	Parallel continuous-time evolution removes the need for algorithmic iteration [8]

4.7 Scalability and Extensibility

Because each cell computes autonomously, the solution time does not increase with the number of discretization points; only the analog bandwidth and settling time limit scaling. The architecture can be extended to three dimensions by stacking dies with through-silicon vias (TSVs) and using vertical coupling MOSFETs, enabling volumetric PDE simulations. Moreover, the same hardware can be repurposed for other coupled-map lattices (e.g., cellular neural networks) simply by reprogramming the diode slopes and coupling conductances.

5. Results

5.1 Single Cell Dynamics

Transient SPICE simulations of the standalone unit cell validated several things such as stable double-scroll attractors, correct breakpoint slopes for the Chua diode, oscillation at supply voltages as low as 0.2 V, and power consumption in the microwatt regime. These results show that the nonlinear reaction term f(u) can be implemented in CMOS with compact DTMOS-based structures.

Transient SPICE simulations were performed on the standalone CMOS unit cell in order to validate its dynamical and electrical behavior. The implemented piecewise-linear Chua diode reproduced the nonlinear dynamics needed for chaotic oscillation since the circuit portrayed stable double-scroll attractors in the phase space. The breakpoint slopes of the I–V characteristic lined up with the designed inner and outer slope values, confirming the correct realization of the programmable reaction term f (V). Sustained oscillation was seen at supply voltages as low as 0.2 V, showing proper subthreshold DTMOS operation and verifying the low-voltage design assumptions. Measured directly from supply current in SPICE transient analysis, average power consumption was still in the microwatt regime, consistent with inductor-less CMOS Chua implementations [1,9]. As a whole, these conclusions highlight that the proposed CMOS architecture implements the intended nonlinear reaction dynamics with low-power and low-voltage operation that's enough for array integration.

5.2 Verification of Diffusion Coupling

2 cells connected using a triode-biased MOSFET were simulated in order to measure diffusion effects. As we predicted, convergence rate was proportional to the programmed gate voltage, node voltages converged toward each other exponentially, and the observed effective g_m matched the expected value from the MOSFET models. This goes to show that the MOS-triode coupling network implements the discrete Laplacian operator accurately.

5.3 Reaction-Diffusion Behavior in an 8x8 Array

An 8×8 array was co-simulated by utilizing a Verilog-A behavioral model. With random initial conditions, the array showed spontaneous symmetry breaking, wavefront propagation behavior consistent with reaction–diffusion theory, formation of stationary spatial patterns (akin to Turing patterns), and steady-state structures matching finite-difference PDE simulations. Convergence was measured in simulation as the time needed for node voltages in order to reach steady-state pattern formation, which happened within around 30–40 nanoseconds of analog simulation time (nominal parameters).

5.4 Impact of Diffusion Tuning

Sweeps of the coupling gate voltage showed that de-

creasing g_m did slow down pattern dynamics and broadened spatial modes, increasing g_m caused faster pattern formation and shorter wavelength structures, and the overall system continued to remain stable across the full tuning range. This tunability supports real-time reconfiguration of the PDE parameters.

5.5 PVT Robustness

Monte Carlo simulations incorporating transistor mismatches, threshold variations, and temperature shifts showed that spatial patterns persisted with small wavelength variations, chaotic attractors maintained structural integrity, and diffusion coupling remained monotonic and well-behaved. So, these simulations show tolerance to fabrication variability.

6. Conclusion

In addition to demonstrating the potential of using coupled chaotic oscillators as a computing substrate to solve nonlinear partial differential equations (PDEs) in hardware, this work outlines how substantial increases in computational speed and decreases in power consumption could be achieved in principle by translating mathematical operations into analog electrical circuits and combining these into a single CMOS chip. The proposed framework further illustrates how complex spatiotemporal patterns could emerge from localized coupling within chaotic oscillator arrays. In contrast to other approaches which used large amounts of space for their components and required milliwatts of power, this new method enables the possibility of integrating many analog processing elements onto a small piece of silicon. Future work will call for fabrication of a prototype array in order to validate the theoretical and simulation results presented here experimentally and to also examine power efficiency and real-world scalability.

Access references for this paper using the QR code below.

Aggression After Traumatic Brain Injury: A Systematic Review of Prevalence, Mechanisms, and Outcomes

Paige Pieterick

Ms. Paige Pieterick is an undergraduate student at The Georgia Institute of Technology double majoring in Biomedical Engineering and Spanish with a minor in Neuroscience. Her current research focus at the Emory School of Medicine involves treatment options for osteodegenerative diseases such as Arthritis and Osteoporosis and focuses on craniofacial abnormalities. She has an interest in neuroscience research and hopes to pursue a MD/PhD focusing on neurological applications.

Abstract

Traumatic Brain Injury (TBI) is a leading cause of long-term behavioral dysfunction worldwide. Aggression represents one of the most disruptive outcomes of TBI lasting well after the acute recovery phase and imposes disruptions to daily life. With limited treatment options, it imposes profound sociofamilial costs. This review examines the prevalence and presentation of aggression following TBI, delineates the underlying neurobiological mechanistic causes of symptoms, evaluates available therapies, and discusses the effects on familial burdens. Aggression arises in 11-34% of individuals post-TBI with verbal aggression predominating (Rao et al. 2009). Imaging studies associated aggression post-TBI with changes in white matter integrity in the corpus callosum and thinning of the orbital frontal cortex (OFC). Lesion studies similarly found frontal and orbitofrontal damage correlating with higher scores on Anger and Hostility scales. Psychiatric comorbidities exacerbate serotonergic and dopaminergic imbalances amidst previous diagnosis of Major Depressive Disorder as a leading risk factor for development of aggression.

Changes in daily life and familial stressors sustain aggression, leading to familial conflict and caregiver burnout with over 70% of caregivers reporting moderate to severe personal distress, correlating with frequency and intensity of aggressive episodes (Ergh et al. 2002). Treatment options vary from known pharmaceuticals for aggression in non-TBI individuals to psychological approaches. Selective serotonin reuptake inhibitors (SSRIs) and Beta-adrenergic antagonists suppress impulsivity and hyperactivity that is associated with reactive aggression; however, limited individualization of treatment results in limited success and caregiver burden as an unaddressed consequence of aggression. Complex neural disruptions, neurotransmitter dysregulation, and physiological discourses drive aggression following TBI. Understanding these factors is essential in treating individuals not only with TBI but other brain injuries resulting in aggression.

Introduction

Aggression constitutes one of the most distressing consequences of TBI and affects not only the individual's treatment, but also familial interactions. The Centers for Disease Control and Prevention estimate that approximately 1.7 million Americans sustain a TBI annually, with mild TBI accounting for around 80% of cases (Faul et al. 2010). While behavioral changes after TBI have been reported, aggression is a distinct change characterized by intentional harm-related behaviors (verbal or physical) that targets others or their environment. These behaviors are socially disruptive but are also indicative of longer hospitalizations and increased risk of readmission (Kim and Bijlani 2006).

Irrespective of the prevalence of TBI and aggression, research and treatment often focus on motor deficits, leaving the production of aggression under researched and treated as a behavioral complication rather than its own syndrome. Imaging studies show decreased white matter integrity in the corpus callosum that correlates with post-TBI aggression. Treatment options do not consider the contribution of neural mechanism dysfunction or caregiver burnout that directly affects rehabilitation outcomes. With no place in treatment options, families suffer, often becoming victims of aggression, which creates a hostile environment that can lengthen recovery time. A lack of research on treatment options for individuals with a history of TBI further exacerbates treatment gaps. The aforementioned factors suggest that aggressive behavior post-injury should be considered as a neural disruption rather than a TBI symptom. The following sections address neural, neurochemical, and social findings of aggression as an effect of disrupted neural circuits and feedback loops following brain trauma.

Review Methodology

To identify primary literature examining aggression following traumatic brain injury (TBI), a literature search was conducted across multiple databases, including PubMed and Google Scholar. Searches for peer-reviewed primary studies published between January 1990 and December 2025 were identified using combinations related to TBI and aggression. Studies involving human participants with a history of mild, moderate, or severe TBI that assessed aggression as a primary or secondary outcome were included. Review articles, case reports with fewer than five participants, and studies focusing exclusively on non-aggressive neuropsychiatric outcomes were excluded. The initial search identified 35 articles, of which 29 met inclusion criteria and were included in the final review.

Traumatic Brain Injury

Traumatic brain injury is an alteration of brain function leading to temporary or permanent cognitive, physical, and behavioral changes (Ginsburg and Smith 2025). Common etiologies of TBI are falls, motor vehicle collisions (MVC), physical assault, and sports-related injuries. The Glasgow Coma Scale (GCS) evaluates the cognitive function of an individual to determine TBI severity. A mild score ranges from 13 to 15, a moderate injury ranges from 9 to 12, and severe impairments fall below nine. Imaging post-TBI revealed diffuse axonal injury (DAI), hemorrhage, and skull fractures incite TBI symptoms. Mechanical forces cause TBIs that act along the white-matter of the brain, causing the shearing of axons. These forces disturb the excitatory and inhibitory neural networks that balance emotional regulation. Acute symptoms following injury include confusion, amnesia (memory loss), and loss of consciousness. Chronic effects include cognitive and personality changes including emotional dysregulation. Rehabilitation efforts include pharmacologic treatments and long-term cognitive therapies (National Institute of Neurological Disorders and Stroke 2025).

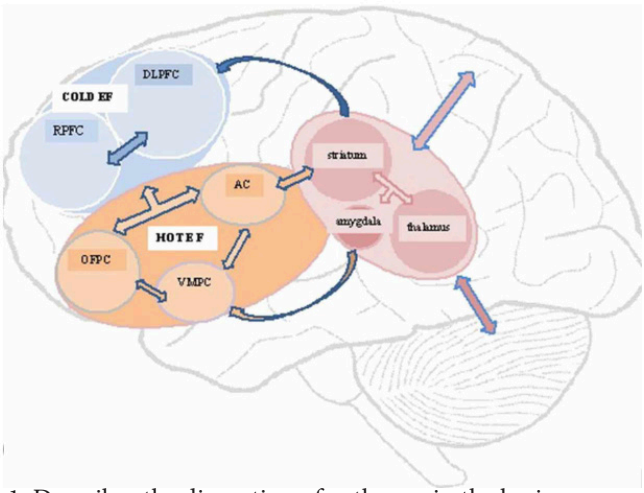

Figure 1. Describes the disruption of pathways in the brain contributing to common TBI-like symptoms and effects (Wood and Worthington 2017).

As seen in Figure 1, common etiologies of TBI symptoms including depression, aggression, and cognitive impairment are caused by damage to the areas of the OFC, thalamus, and amygdala specifically. Damage to the OFC reduces inhibition of other

cortical areas while the thalamus and amygdala affect emotion regulation through the fear and aggression response of the Amygdala (Wood and Worthington 2017). Disruptions to these areas implicate symptoms of aggression, depression, and anxiety because of excitation of the amygdala and thalamus.

Aggression

Aggression following TBI is one of the most disruptive symptoms of a brain injury and causes mental health effects and social dynamic changes. Verbal aggression including yelling, swearing, or threatening, is the most common type of aggression following a TBI; however, physical aggression, property destruction, and self-directed harm can also occur (Rao et al. 2009). Aggression after TBI is categorized into reactive and instrumental types. Reactive aggression is consistent with impulsivity and is emotionally driven as a response to perceived threats. For example, throwing objects in response to aversive sensory stimuli. Whereas instrumental aggression is a goal-directed subtype of aggression focused on personal gain. For instance, a patient refusing (hostilely) to abide by treatments unless a certain demand is met. Reactive aggression predominates after TBI injury and aligns with deficits seen in emotional regulation from the Amygdala. Instrumental aggression is not seen as often and is likely due to a lack of decreased executive function. Identifying the subtype of aggression at hand is a necessity because pharmacologics that target impulsivity (Reactive Aggression) may not work for those with goal-directed aggression.

Aggressive patients also demonstrate an increased risk of substance abuse, non- compliance to treatments, and difficulties reintegrating into society. Socially, aggression imposes strain on relationships and diminishes patient quality of life, leading to involvement in illegal activities due to decreased activation of dopaminergic reward pathways. Caregivers of these patients report a heightened stress level and caregiver burnout as a direct result of these behaviors (Kolakowsky-Hayner, Miner, and Kreutzer 2001). Because aggression is often classified as a behavior, it receives limited specific study in the presence of TBI, resulting in generalized treatments instead of ones tailored to the individual. Understanding aggression mechanisms is critical for improving both patient and family outcomes by individualizing care.

Mechanisms of Aggression

Neural Mechanisms

Neuroimaging post-TBI consistently demonstrates that there are disruptions of the limbic system common in patients with aggressive tendencies. Patients commonly diagnosed with a mild TBI show Elevated radial diffusivity (a measure of White Matter damage) in the Corpus Callosum (Dailey et al. 2018). These findings also correlate with total aggression scores on the Buss-Perry Aggression Questionnaire, which evaluates the patient on the following symptoms of the four aggression subtypes: physical aggression, verbal aggression, anger, and hostility. A higher score on the questionnaire correlates with decreased white matter integrity within the splenium and callosum body.

Buss-Perry Aggression Questionnaire (BPAQ)

Instructions:
Using the 5 point scale below, indicate how uncharacteristic or characteristic each of the following statements is in describing you.

		Extremely Uncharacteristic	Somewhat Uncharacteristic	Neither Uncharacteristic Nor Characteristic	Somewhat Characteristic	Extremely Characteristic
1	Some of my friends think I am a hothead.	1	2	3	4	5
2	If I have to resort to violence to protect my rights, I will.	1	2	3	4	5
3	When people are especially nice to me, I wonder what they want.	1	2	3	4	5
4	I tell my friends openly when I disagree with them.	1	2	3	4	5
5	I have become so mad that I have broken things.	1	2	3	4	5
6	I can't help getting into arguments when people disagree with me.	1	2	3	4	5
7	I wonder why sometimes I feel so bitter about things.	1	2	3	4	5
8	Once in a while, I can't control the urge to strike another person.	1	2	3	4	5
9	I am an even-tempered person.	5	4	3	2	1
10	I am suspicious of overly friendly strangers.	1	2	3	4	5
11	I have threatened people I know.	1	2	3	4	5
12	I flare up quickly but get over it quickly.	1	2	3	4	5
13	Given enough provocation, I may hit another person.	1	2	3	4	5
14	When people annoy me, I may tell them what I think of them.	1	2	3	4	5
15	I am sometimes eaten up with jealousy.	1	2	3	4	5
16	I can think of no good reason for ever hitting a person.	5	4	3	2	1

Figure 2. Shows the Buss-Perry Aggression Questionnaire with characteristics for verbal, physical, generalized anger, and hostility aggression. This questionnaire aids in diagnosing the specific type of aggression present (Buss and Perry 1992).

Structural MRI findings showed thinning of the left lateral OFC, associated with irritability and impulsive behaviors (Epstein et al. 2016). Lesion studies in veterans also identified frontal and orbitofrontal damage correlation with higher scores on the Anger–Hostility scale through the Buss-Perry Aggression Correlate (Grafman et al. 1996), and patients with frontal lobe lesions exhibited a significantly higher frequency of aggression than those with diffuse axonal injuries (Tateno, Jorge, and Robinson 2003).

These findings imply that there are disconnections between the amygdala and the OFC, which are structural pieces of the aggression pathway and are commonly disrupted in TBI. The OFC modulates inputs from the amygdala to allow for re-evaluation of aggression and context-dependent response. Damage to this feedback circuit is important as

individuals show poor risk assessment, risk evaluation, and impaired decision making in conjunction with damages to the areas illustrated in Figure 3.

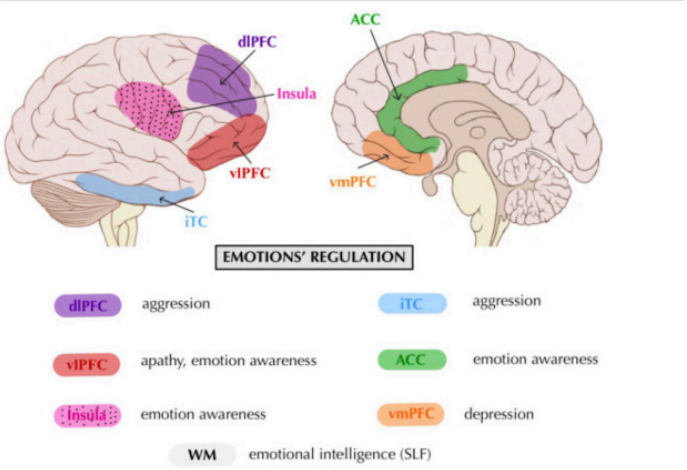

Figure 3.Shows the regions involved in emotion regulation and aggression following TBI. There are connections between the dorsolateral prefrontal cortex (dlPFC), ventrolateral prefrontal cortex (vlPFC), and inferior temporal cortex (iTC) which are associated with aggression and lack of control. The Insula and Anterior Cingulate Cortex (ACC) mediate emotional awareness and empathy. The ventromedial prefrontal cortex (vmPFC) contributes to mood regulation and negative affect (Cristofori et al. 2024).

Damage to the prefrontal cortex, particularly the dlPFC and vlPFC, results in impulsivity and reactive aggression due to their roles in decision making and logic. Anterior cingulate cortex (ACC) lesions impair empathy and social cues, while those to the insula contribute to irritability (Cristofori et al. 2024). Injuries involving the vmPFC also cause dysregulations and show comorbidities with depression and anxiety caused by mood instability. Decreased white matter integrity within the superior longitudinal fasciculus (SLF) and uncinate fasciculus creates a miscommunication between understanding our emotions and experience (Cristofori et al. 2024). These patterns explain why aggression after TBI is often attributed to blunting of emotions, apathy, and a depressive affect.
Together, these studies demonstrate that post-TBI aggression is not attributed to a single neural dysfunction but instead affects many intertwined networks. This shows that aggression is a result of reduced feedback from the OFC to structures like the amygdala. By showing these interactions, Figure 3 shows that successful treatment could also be found by looking at communication systems. Treatment options should take these damages into account as they affect emotional regulation, producing impulsive and reactive aggression (Cristofori et al. 2024).

Neurochemical and Psychiatric Comorbidities
Aggression after TBI is also influenced by neurochemical imbalances that further amplify structural dysfunctions. Decreased levels of serotonin evident by the reduction of cerebrospinal 5-HT levels, has been associated with reactive aggression and impulsivity (Cantor et al. 2005). Dopaminergic dysregulation could impair reward centers and cause a lack of behavioral inhibitions (Ruffieux et al. 2018). Chronic GABAergic (Gamma-Aminobutyric Acid Synapses) deficits further affect mood regulation because of feedback loops (Marcelo Liborio Schwarzbold et al. 2010). Psychiatric comorbidities exacerbate these neurochemical imbalances and are commonly diagnosed six months post-TBI. Major depressive disorder and anxiety disorders correlate with irritability and anger and further reduce the ability to suppress impulsive aggression (Rao et al. 2009).
The interaction between neurotransmitter imbalance, structural dysfunction, and psychiatric comorbidity explains persistent aggression after the initial phase of recovery. Considering these interactions is necessary in crafting interdisciplinary treatments that navigate multiple avenues of potential causes.

Prevalence of Aggression Post-TBI

Methods
Researchers employ various methods for post-TBI behavioral study. One of these assessments, the Overt Aggression Scale (OAS) quantifies the frequency and severity of self-reported symptoms regarding verbal aggression, physical aggression, and property destruction, while the Neuropsychiatric Inventory (NPI) and Behaviour of Concern Scale (BoC) use information provided by family and caregivers. Because of this, discrepancies may arise given the lack of consistent measurements as self-reported information can be inaccurate. Differences in measurement spanning across different time periods post-injury also affect results and outcomes. Acute studies often have findings of irritability and agitation which may not be signs of true aggression. While later, follow-up testing assesses chronic patterns of aggressive behaviors determined to be consistent with a true change in behavior. Sample sizes are also small, and most exclude individuals with pre-existing psychiatric illnesses. Inconsistent sample

sizing and a lack of scheduled follow-up testing limit the reliability of prevalence estimates.
Aggression prevalence post-TBI ranges from 11% to over 40%, depending on the timing of the study after TBI, and the TBI severity determined by the GCS scale. During the first three months of injury, aggression is present in 28.4% of TBI patients with verbal aggression, depression symptoms, and dependence on caregivers prevailing (Rao et al. 2009).Persistent aggression when measured between 6-12 months post-TBI was reported in 38-41% of patients, with males and those of lower education levels showing higher levels during this period (Roy et al. 2017). These persistent behavioral changes were confirmed in a study conducted five years post-TBI resulting in a 25% prevalence rate of clinically diagnosed aggression (Wade, Michaud, and Brown 2006), with 70.5% of patients reporting at least one behavioral change, verbal aggression predominating (Hicks et al. 2017). These behavioral changes correlate with increased familial stress that prolongs symptoms and increases risk of non-compliance to treatment.

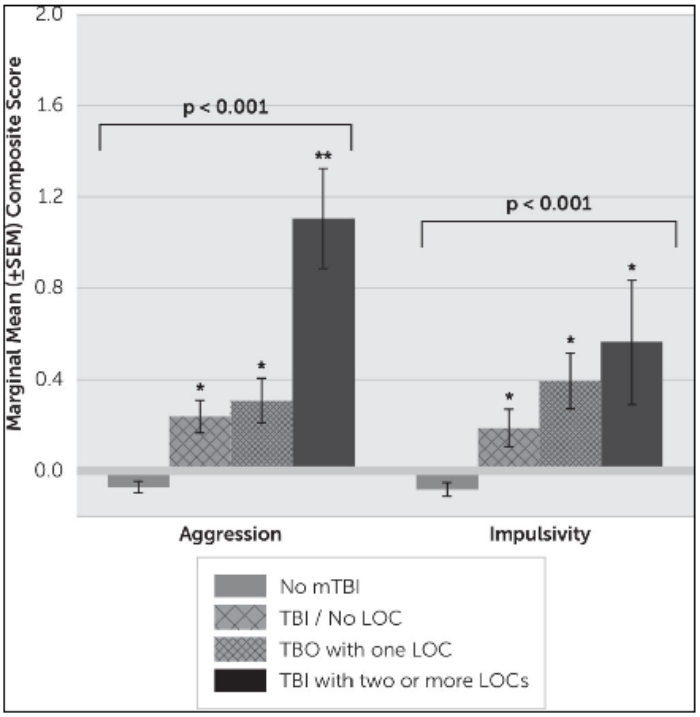

Figure 4. Describes the mean aggression and impulsivity scores correlated with TBI severity and loss of consciousness (LOC) frequency. Individuals with multiple TBIs and two or more LOCs demonstrated significantly higher aggression and impulsivity compared with all other groups ($p < 0.001$). Error bars represent ±SEM (Roy et al. 2017).

Figure four demonstrates that increased severity of TBI, often characterized by increased rate of loss of consciousness, is a determinant in the prevalence of aggression in patients. These findings suggest that acute irritability, recognized in initial aggression studies (3 months post-TBI), frequently develops into chronic aggression. This transition is characterized by a period of reactivity followed by personality and mood changes consistent with true aggression.

Time Post-TBI	Prevalence	Aggression Type	Correlate
3 months	28%	Verbal	Depression
6–12 months	38–4%	Verbal	Male sex, low education
>1 year	70%	Verbal/social	Familial stress
5 years	25%	Verbal	TBI severity

Table 1. Describes a brief overview of the aforementioned studies examining prevalence of aggression across different time spans throughout the recovery period (Rao et al. 2009; Marcelo Liborio Schwarzbold et al. 2010; Wade, Michaud, and Brown 2006; Hicks et al. 2017).

Current Treatments

Clinicians most commonly employ Pharmacotherapies; however, they are the least individualized method of treatment. Selective Serotonin Reuptake Inhibitors (SSRIs), such as Sertraline (Zoloft) and Citalopram (Celexa) increase levels of serotonin in the synapse, and have been shown to mitigate impulsivity (Neurobehavioral Guidelines Working Group et al. 2006). Beta-Adrenergic antagonists like Propranolol, inhibit the sympathetic nervous system and reduce hyperactive symptoms (CASPER 1991). Antipsychotics, for instance Quetiapine and Olanzapine, have had success in controlling aggression but also have serious side effects like sedation and metabolism irregularities. Additionally, mood stabilizers such as Valproate and Carbamazepine are understudied but have been promising in reducing aggression (Kim 2002).
Pharmacological approaches appear promising; however, aggression testing of pharmaceutical drugs has not been completed in individuals with a history of brain injuries; thus, lacking the knowledge of appropriate dosing, potential side effects for those with structural dysfunctions, and efficacy. Medications also target a single neurotransmitter which may not be effective as evidence supports aggression post-TBI as the result of multi-system dysfunction affecting the

limbic systems. Additionally, serious side effects of antipsychotics also impede on patient quality of life and run the risk of non-compliance to medication. Cognitive Behavioral Therapy (CBT) used for TBI focuses on identifying anger and provides skills for relaxation and trigger identification (Ponsford et al. 2003). Programs focused on social skills use techniques that facilitate social cues and implement deescalation strategies for family therapies. These family-based therapies acknowledge that caregiver distress directly influences patient outcomes.

Combination therapies that utilize both the pharmaceutical and cognitive approaches work by pairing SSRIs with behavioral interventions to target the biological and psychological disruptions. However, trials of these techniques have rarely been employed and have only been tested on specific aggression types (verbal). In response to this, recent techniques including transcranial magnetic stimulation (TMS) and neurofeedback have shown promise in reinstating inhibitory control of the frontal lobe (Rennert et al. 2016). Despite this, current treatment protocols still neglect the individualized approach towards treatment plans that can be tailored to patient needs. Using imaging markers like white-matter integrity and OFC thickness, pharmaceuticals and behavioral therapies could be tailored to the individual and would predict a higher rehabilitative response.

Treatment Type	Mechanism / Approach	Limitations
Pharmacologic	SSRIs (sertraline, citalopram), β-blockers (propranolol), mood stabilizers (valproate, carbamazepine), and antipsychotics (quetiapine, olanzapine).	Reactive aggression. Risk of sedation and dizziness.
Psychological / Behavioral	Cognitive-behavioral therapy (CBT) and anger management training. Focuses on Social Cues	Used for mild to moderate TBI. Lack of available therapists.
Family-Centered	Family therapy and caregiver support programs reduce environmental triggers and improve communication; Could use TMS and neurofeedback.	Better caregiver outcome and patient adherence to the treatment plan. Costly and under-studied.

Table 2. Provides a summary of the most commonly employed interventions and their categories for post-TBI aggression and also explains the integration of a complex treatment plan with approaches that implement behavioral tools and family-system approaches.

Familial Effects

Aggression after TBI impacts family systems and caregivers of those with aggressive tendencies. Over 70% of caregivers report moderate to severe stress when caring for a family member with aggression (Kolakowsky-Hayner, Miner, and Kreutzer 2001). Caregiver burden also directly correlates with the number of aggressive episodes (Ergh et al. 2002). Reactive aggression creates a hostile environment in the home, leading to burnout as episodes become more prominent and higher hostility scores in patients predict greater caregiver burden even when taking injury severity into account (Klyce et al. 2023). These results demonstrate that burnout is independent of physical limitations from the injury but is rather affected by the environment in the home. The Transactional Stress–Coping Model (Lazarus and Folkman 1984), describes stress as a culmination of situations and available resources and confirms these results. Spousal caregivers, often women, exhibit high anxiety due to the roles as partner and caregiver at the same time (Marsh et al. 2002).

The consequences of this stress disrupt work, social life, and financial life balances. This cycle impacts the environment at home and further exacerbates patient stress, decreasing quality of life. This dynamic is critical to patient outcomes. Despite this, healthcare systems provide little caregiver support in the realm of aggression treatment. Programs involving emotional regulation skills reduce stress and improve family dynamics (Shepherd-Banigan et al. 2018); however, these programs are seldom implemented or included in treatment plans. Approaches that address the neural disruptions, neurotransmitter imbalance, and home environment, should be implemented when attempting to treat patients with brain injury. Focusing on mitigating environmental triggers improves family life and creates an environment conducive to rehabilitation.

Conclusion

Aggression after TBI is a complex syndrome spanning across neurological mechanistic interruptions and chemical disruptions that are influenced by the environment. Across studies, disruptions in the OFC and amygdala and reduced white matter integrity along with serotonergic and dopaminergic disruptions lead to impulsivity and irritability consistent with reactive aggression. Individuals with pre-existing psychiatric conditions, recurrent brain injuries, and severe impairment disproportionately develop long-term behavioral changes when irritation is present within the first three months. This suggests that care prioritizing motor deficits may disregard behavioral symptoms, leading to potential consequences that affect quality of life.

Thus, aggression should be referred to as a dysfunction rather than a secondary behavioral symptom. However, methodological inconsistencies and limited studies in post-TBI aggression specifically, prevent improvement of the current diagnostic system. The variability in prevalence study results reflect methodological limitations including exclusion of individuals with psychiatric comorbidities and the lack of standardized follow-up procedures. Utilizing assessment tools that combine the methods of the Buss-Perry Scale, OAS (self-reporting of symptoms), BoC, and NPI (family reporting of symptoms) standardizes diagnostic tools and allows for accurate diagnosis.

Treatment options are further limited by standardization that does not consider a multi-disciplinary approach. Pharmacological interventions disregard the evidence that supports aggression as a multi-level dysfunction, and behavioral therapies are not enough to overcome neural dysfunction. Approaches that incorporate neuroimaging into care could indicate disruptions in the amygdala, dlPFC, and white matter integrity of the corpus callosum that are helpful in determining the treatment method necessary (SSRIs, Beta-Blockers, or Antipsychotics). Novel approaches like TMS and Neurofeedback, when incorporated with family-centered therapies that positively affect the home environment, could aid in patient outcomes with continued research as preliminary results are promising.

The combination of psychological predisposition and environmental stressors should also be taken into account when considering treatment options since neural mechanism dysregulations are key causes of aggression. Even with this in mind, environmental effects remain under-researched. Because caregiver burden predicts success rates of rehabilitation, family treatment should be integrated into standard TBI treatment. However, support for caregivers is inconsistent and is rarely considered. Effective treatments should consider the environment in which recovery occurs and the best way to create a positive family dynamic. Future advancements should include both the patient and the caregiver and at the same time, address aggression as an indicator of neural mechanism dysfunction rather than a behavioral issue. Addressing this requires a shift in thinking, inclusion of multi-modal diagnostic techniques, and individualized treatment options that support patient quality of life. This represents a key step toward restoring individuals' daily routines back into their lives post-TBI.

Acknowledgements

I gratefully acknowledge Dr. Jasmine Hope and Dr. Mary Holder for their careful review of the manuscript and for providing valuable feedback that greatly strengthened the final paper. I would also like to thank Dr. Steven Goudy and Dr. Hope Robinson for their guidance and support throughout the research process. Their mentorship has been essential in my personal growth as an aspiring scientist.

Access references for this paper using the QR code below





JOIN OUR TEAM

OPEN POSITIONS:

Assistant Manager (Operations/Maintenance) • **Crane Mechanic** • **Crane Operator**
EMTs (Electronic Mechanical Technicians) • **Communications and Purchasing**
Equipment Operator • **Information Technology** • **Marketing** • **Port Police**

THREE LOCATIONS:

Savannah • **Brunswick** • **Crandall**

EXCEPTIONAL BENEFITS PACKAGE:

Affordable Medical/Dental/Vision Insurance — Starting on Day 1! • Wellness Program
Company-paid Life Insurance • Matching Retirement Plan • 12 Paid Holidays/Year • Yearly Company Bonus
Tuition Reimbursement • Generous Vacation/Sick Leave • On-site Medical Clinic

GEORGIA PORTS

SCAN THIS QR CODE

Explore the possibilities at gaports.com/jobs

All positions must meet the criteria for and obtain a Transportation Worker Identification Credential (TWIC). Valid state driver's license required.



Closing the Coverage Gap: Reforming Georgia's 1332 Waiver to Advance Rural Health Insurance Equity

Calvin Zhu, Sakthi Priya Ramamoorthy, Sri Dhanya Muppalla, Ria Vittal

Calvin Zhu, Sakthi Priya Ramamoorthy, Sri Dhanya Muppalla, and Ria Vittal are all members of the Health and Biotechnology Committee within the organization of Science and Technology Policy Connections (S&TPC) at Georgia Tech. Through S&TPC, students craft innovative policy solutions to prominent issues, empowering students to shape science and technology policy for a safer future. The authors are specifically interested in promoting healthcare equity through healthcare policy and public health initiatives. They hope to improve accessibility by ensuring widespread access to high quality healthcare.

Abstract

Many Georgians living in rural communities suffer from a lack of accessible healthcare as a result of soaring premiums from insurance agencies that market "affordable" plans with high deductibles and delayed treatment. There is a need for governmental intervention to increase the number of insured individuals while ensuring equity in the cost of premiums and quality of care. However, the Georgia Access Model, in its current state, fails to guarantee affordable and comprehensive coverage. Georgia's Section 1332 waiver consists of a two-part approach to address healthcare access and affordability challenges: (1) a state reinsurance program and (2) the Georgia Access Model, a major overhaul of the individual health insurance market. While the first phase is standard, the second phase introduces policies that undermine the Affordable Care Act (ACA) and are legally questionable. The Georgia Access Model, a state-run marketplace for insurance, will enable the addition of insurance plans that are neither affordable nor comprehensive, which will limit the already scarce options that rural Georgians have when it comes to healthcare. With further budget cuts, Georgia's Section 1332 Waiver would fail to meet its goals of expanding coverage for vulnerable citizens, particularly those living in rural communities.

.

Introduction

Georgia's lack of affordable healthcare coverage is a crisis that disproportionately impacts rural communities. Approximately 55 rural counties have a higher rate of uninsured residents compared to the state average, which is already the second-highest uninsured rate in the country [14]. Even for those who are insured, high medical costs prevent 15.2% of people from seeing a doctor when they need to. These barriers have put Georgia among the top five worst states for healthcare access, leaving thousands without essential care [28].

Georgia's Section 1332 waiver uses a two-part strategy to address affordability and coverage issues. The waiver introduces a state reinsurance program and, more significantly, the Georgia Access Model, which is a fundamental restructuring of the individual insurance marketplace that shifts enrollment away from the federal exchange and towards private insurers and brokers. This shift has made it even more difficult for low-income and rural residents to find affordable plans. Instead of improving access, this approach has led to fewer consumer protections and increased confusion about coverage options.

Georgia's healthcare crisis is already worsening under the state's 1332 waiver and the Georgia Access Model, which has only deepened existing coverage gaps instead of closing them. Under the expiring Affordable Care Act (ACA) tax credits, many families would face premiums and out-of-pocket costs so high that coverage becomes a luxury. Estimates from the Congressional Budget Office show that the "One Big Beautiful Bill" Act would lead to an estimated "560,000 - 900,000 Georgians" and over "10 million people nationwide" losing healthcare insurance by 2034 [9]. In states like Georgia, where rural residents already face low insurer competition, sparse risk pools, longer travel distances for care, and higher baseline premiums, these subsidies are instrumental in keeping people from plunging into uninsurance.

To mitigate these harmful effects, Georgia state legislators should consider potential policy solutions such as expanding Medicaid, strengthening outreach efforts, and increasing oversight on brokers. Without decisive action, healthcare will be pushed further out of reach for the millions of Georgians who need it most.

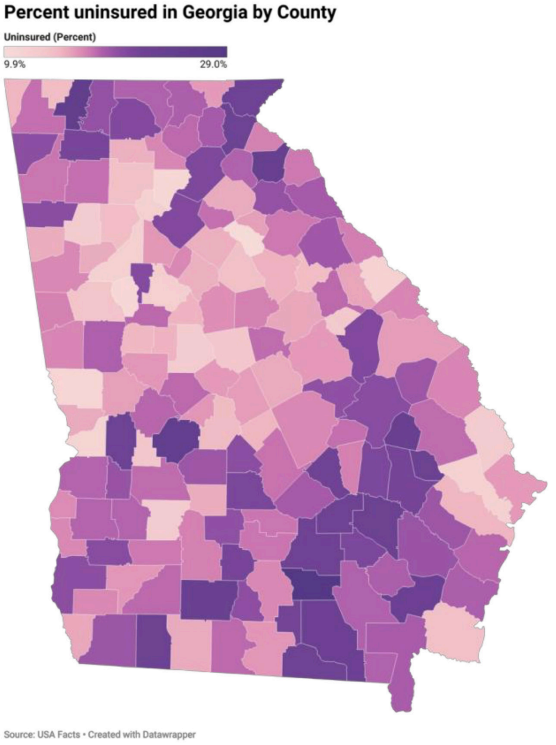

Figure 1. Percent uninsured across all income levels for individuals age 18-64 in Georgia by County in 2023. Data was derived from HDPulse [25]

.

We aim to evaluate how the Georgia Access Model has altered insurance coverage and consumer protections for rural Georgians and examine which policy interventions are the most feasible to mitigate coverage loss. This paper examines Georgia's Section 1332 waiver with a focus on healthcare access in rural communities and proposes policy recommendations to mitigate healthcare inequities and coverage. Through a policy analysis and evaluations of feasibility, we aim to provide evidence-based guidelines for policymakers and community stakeholders.

Repercussions for Georgia

The implementation of the Georgia Access Model, in its current state, cuts off federal funding for marketing and outreach to rural communities. Georgia, if unwilling to compensate for this loss of funds, will see a drop in non-group enrollment, specifically among rural Georgians. Additionally, moving away from the Healthcare.gov platform increases attrition, as the removal of the federal marketplace will inevitably

increase confusion among enrollees, many of whom are accustomed to the federal marketplace and will leave the market during the transition. Furthermore, Medicaid enrollment, which is important for rural hospital funding and health coverage for uninsured individuals, will be negatively impacted [1]. A decrease in non-group enrollment in the federal marketplace will reduce the number of consumers who are notified of their Medicaid-eligibility status during enrollment, and brokers and insurers, who do not receive commission for Medicaid enrollment, will be disincentivized to enroll eligible consumers in Medicaid [29]. A model of

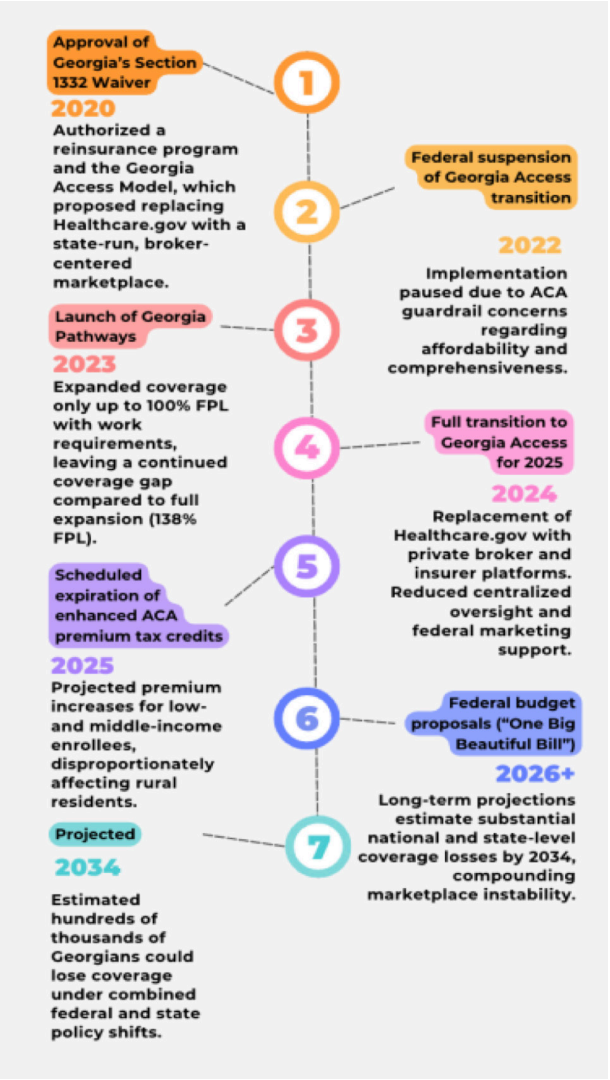

Figure 2. Chronology of the federal and state policy events that affect Georgia's insurance marketplace structure, levels of government insurance subsidies, and long term projections of coverage. The Georgia Access Model implemented led to major structural shifts, replacing the federal marketplace.

2023 enrollment, which considered the three different scenarios adjusted based on the severity of changes in advertising, attrition, and Medicaid enrollment if the Georgia Access Model were to be implemented, projected a consistent drop in Medicaid enrollment across all three scenarios1.

The Georgia Access Model would further this issue by reducing competition among marketplace insurers. To participate in the Georgia Access Model, insurers have to guarantee privacy, security, and plan display standards, which can be a barrier to entry for small insurers [2]. Small insurers will be unable to compete with larger insurers in terms of marketing budget, commission for brokers, and brand recognition. Dominant insurers, who are pressured to lower their plans by other insurers in the federal marketplace, will be able to exercise an unjust amount of power in controlling marketwide premium levels, and notably, larger insurance companies are less likely to sell in rural communities because of smaller risk pools and higher costs, which will be detrimental for insurance access.

Additionally, the current federal landscape has been marked by an impending expiration of Premium Tax Credits under the Affordable Care Act Marketplaces. The Urban Institute Project projects that in 2026, the average net premiums would increase by more than four times for Marketplace enrollees with an income 250% below the federal poverty line (FPL), more than double for those with incomes 250-400% above the FPL and nearly double for higher income enrollees. The loss of such assistance would exacerbate the enrollment declines.

Even as premiums are expected to rise under the Georgia Access Model, the quality of care offered by these plans is expected to worsen. The waiver allows subsidies for short-term, limited-duration insurance (STLDI) and association health plans (AHPs), which will lead to less comprehensive coverage for those with preexisting conditions, thereby only encouraging healthy individuals to get insurance. More at-risk populations, such as those in rural areas, will have higher premiums. Furthermore, the state would allow for the sale of "copper plans," which are non-ACA-compliant plans that require significantly higher cost-sharing than is permitted by the ACA. The trade-offs for the lower premiums of copper plans include higher deductibles

and excessive out-of-pocket spending by policyholders, which increases the likelihood that enrollees, particularly in rural communities, will delay or forgo care.

Policy Options

To address the critical issues presented by the Georgia Access Model, we propose three possible policy alternatives that can be implemented independently or in some combination.

i. AI-based Oversight on Brokers

A transition to the Georgia Access Model results in private insurer agents and brokers being the primary means through which people select and purchase their insurance. With limited oversight from the state, web-brokers are allowed to advertise plans that are not ACA compliant, meaning that they lack one of the following essential health benefits: ambulatory patient services, emergency services, hospitalization, maternity and newborn care, mental health and substance use disorder services, prescription drugs, rehabilitation and habilitative services and devices, laboratory services, preventive and wellness services, chronic disease management, and pediatric services. The Georgia Access Model recently released their 2025 Open Enrollment findings and has mentioned that 562 private agents had their GA Access Model Certification terminated due to improper or fraudulent practices [15]. Given the high proportion of fraudulent agents and brokers, more stringent oversight is needed to identify deceptive activity before people are affected. Moreover, directing people to purchase these plans will have a disproportionately negative impact on rural Georgians, who are already burdened by travel distances and hospital closures.

While an Interactive Virtual Agent (IVA) already exists within the framework of Georgia Access, its capabilities are limited to automated assistance for routine issues, such as navigation. The implementation of a more tailored AI agent could address concerns by bolstering broker oversight and consumer protection via plan recommendations. With approval from CMS, the AI agent can be trained on marketplace plan data and deidentified application outcomes. Three categories of data are needed for training of an AI agent: (1) consumer data, including anonymized household income, age, household size, geographic location, and previously enrolled plans, (2) plan attributes, such as premiums, cost-sharing, and provider networks, and (3) broker activity, including but not limited to the plans recommended to consumers, commission, and final plan selection. Moreover, there are stringent documentation requirements regarding protected health information that must be met in order to comply with the HIPAA Security Rule. The minimum necessary standard set by HIPAA requires that an entity identify the categories or types of protected health information needed [26]. Therefore, types of collected PHI will need to be explicitly defined in order to minimize privacy concerns. Furthermore, data-sharing agreements would need to be reached with CMS, insurers, and brokers before an AI contractor can begin model development. All data would be stored in a secure data environment (SDE) with role-based access and would have to adhere to HIPAA and privacy laws, which can be accomplished by removing any personally identifiable information and concealing broker identities prior to model training.

The AI agent would be evaluated on accuracy, defined by its ability to flag suspicious recommendations that: propose non-ACA compliant plans, provide brokers with abnormally high commission, or exclude subsidized plans. Additionally, to ensure that recommended plans reduce out-of-pocket costs and provide an adequate care network, several metrics will be evaluated, including out-of-pocket cost optimization and plan quality metrics. The median deductible and out-of-pocket maximum from AI-recommended plans will be compared against those of brokers, and the percentage of AI-recommended plans that exceed the ACA out-of-pocket spending threshold should be below that of broker-recommended plans. In terms of plan quality, the percentage of AI-recommended plans that are considered ACA compliant should be higher than that of brokers. Another criterion would be the minimization of the difference in quality of recommendations between urban and rural consumers, which can be determined through the difference in out-of-pocket cost optimization and plan quality metrics between urban and rural counties. After the model has been trained to classify broker behavior and recommend plans, the user interface can be developed with simplicity and plain language in mind, which can be in the form of an interactive chatbot.

Before integration into the current framework, the AI agent would be tested using risk assessments to de-

Table 1. A feasibility matrix for all proposed policy options. All possible policy options were graded on a scale ranging from "1" (low feasibility) to "5" (high feasibility) according to the following six criteria: economic feasibility, technical feasibility, political feasibility, legal feasibility, public support, and efficacy. Total scores were calculated to indicate favorability of policy options.

Criteria	Description	1 Rating	5 Rating	Medicaid Expansion	Strengthen Outreach and Education	AI-based Oversight of Brokers
Economic Feasibility	Is there a sufficient source of funding to implement and sustain the policy? Is the cost-benefit justified for each policy?	No source of funding exists; Not cost effective	Funding exists and can continue to be supplied; Highly cost-effective	4	3	3
Technical Feasibility	Are existing resources able to support the policy? If not, will the policy require a significant investment of time and skills?	No existing infrastructure or resources to support implementation	Resources are ready to be supplied; Infrastructure already exists and can be adapted	4	4	2
Political Feasibility	Is the proposed policy likely to be approved by legislators in the current political environment?	Extremely difficult to be passed by legislators	Easy to be passed by legislators	1	5	4
Legal Feasibility	Does the policy comply with existing state and federal guidelines, and will the policy encounter legal resistance?	Violation of existing laws and guidelines	Compliant with existing laws and guidelines	3	4	4
Public Support	Will the general public view the policy favorably?	Highly negative opinion on policy	Highly positive opinion on policy	4	4	4
Efficacy	Is there a high likelihood of success for the policy?	Coverage and state of rural healthcare is unchanged or worsened	Dramatic improvements in coverage and relief for rural healthcare	4 **Total: 20**	5 **Total: 21**	4 **Total: 21**

termine vulnerability and privacy protection before deployment in randomized counties. Only when key metrics have been met, such as user satisfaction and accurate broker flagging, will the AI agent be employed in a statewide setting through the Georgia Access Model.

Strong oversight is needed to prevent consumers from being taken advantage of and to ensure active efforts are made to increase enrollment. By supplementing AI agents with private insurers and brokers, consumers would feel more reassured and confident, thereby encouraging more engagement on the marketplace and reducing exploitative marketing. This can help establish a more equitable environment for consumers and increase the likelihood that all Georgians receive comprehensive coverage.

ii. Strengthening Outreach and Education

During the shift to the state marketplace, many individuals can be left behind simply because they were not adequately informed. As Georgia switches to the House Budget Proposal of the One Big Beautiful Act in 2026, this issue will be further exacerbated, since the auto-renewal feature is set to be removed. In 2025 alone, around 927,000 Georgians relied on the auto-renewal; however, individuals will now be expected to reapply every year [19]. In their 2025 report, Georgia Access recorded an improvement in the number of insured individuals, but such new changes can reverse this success [15]. To mitigate this risk, health insurance literacy needs to be improved, particularly in rural Georgia. Health education promotion and community outreach will be vital in disseminating information to rural Georgians.

A two-pronged approach is needed to tackle increased attrition and inequitable plans. First, a multimedia campaign involving digital, online, print, radio, and social media will need to be conducted to increase awareness of the shift to Georgia Access. Educational ads, social media campaigns, news segments, community flyers, and radio spots are needed to reach a larger proportion of the rural population in Georgia. In addition to bolstering enrollment, disseminating materials through this multimedia approach will teach people what insurance plans are considered ACA compliant, how to compare plans, and how to recognize substandard plans. Additionally, a "nudging system" consisting of low-cost strategies, such as emails, generalized letters, and automatic phone calls, has previously been shown to increase enrollment in Medicaid [13]. Implementing text-based enrollment efforts can be as simple as using SMS to notify people, including previous enrollees in the federal marketplace, of the shift to Georgia Access, which would be a low-cost and effective method to reach all Georgians.

The other approach to addressing the consequences of Georgia Access would be through community outreach, which has been effective for reaching less informed and previously unenrolled populations. This will consist of partnering with rural institutions, mainly through local clinics, rural hospitals, churches, county health departments, and libraries, in order to reach rural families. Already existing networks can be used to travel to these institutions and deliver information and assistance regarding Georgia Access. For example, Coverage Captains, a program created by Georgians for a Healthy Future, are trusted community members dedicated to improving health care access in their communities. Leveraging this network to increase health literacy would mean partnering with Georgians for a Healthy Future to develop training materials for Coverage Captains and relay information about Georgia Access, including knowledge about how to enroll using the marketplace, the risk of substandard plans, and other resources to improve health insurance literacy.

While Georgia's current approach may still technically support a "no-wrong-door" policy, which claims that applicants will be routed to the coverage best suited for them, there are legitimate concerns about whether consumers are truly ending up with more comprehensive, higher-quality care. This makes transparent, unbiased education all the more essential. Vulnerable rural Georgians should not be targeted with cheap insurance plans that look appealing upfront, but fall short in times of need.

iii. Medicaid Expansion

Currently, Georgia promotes Georgia Pathways to Coverage, which offers Medicaid for people with income less than or equal to 100% of the FPL. It is only a partial expansion of Medicaid, and there are eligibility requirements, including work requirements and premium payments, that suppress meaningful coverage. A complete expansion of Medicaid would raise the

Table 2. This table details the major policy stakeholders involved in the Georgia Access Model, describing their responsibilities, perspectives, and impacts of the policy on each individual group.

Stakeholder	**Analysis**
Georgia Health Policy Center (Advocacy Group)	Instead of using the healthcare.gov platform, the Georgia Access Model aims to use a network of private entities to enroll people in the Affordable Care Act plans. The Georgia Health Policy Center is an organization that collects data and provides evidence-based policy guidance (Georgia Health Policy center, n.d.). They can help facilitate conversations between various stakeholders, are integral to discussions about the Georgia Access Model, and can be involved in monitoring outcomes of the Georgia Access model such as enrollment values and coverage disparities.
Rural Community Organizations	Rural Community Organizations address both local, specific needs, as well broader, more general needs across the nation. These organizations are well equipped and trusted by the members of their community to reach local needs. Specifically, the Navigator and Certified Application Counselor (CAC) program help with outreach efforts as well as ensure that individuals seeking health insurance are well informed and have enrollment assistance (Georgia Access, 2024). This community based approach helps Rural Community Organizations understand specific needs and challenges of a region (Curtis et al., 2015). With the established trust, Rural Community Organizations can prove to be a valuable point of contact for discussions about the impact of the Georgia Access Model as well as how these and other organizations can tailor their efforts to meet the new health insurance model.
Insurance Companies (operating in Georgia)	The products of health insurance companies are at the core of the system, including the changes with the Georgia Access Model. Insurance companies design health plans and charge premiums and copays according to the type of plan and the amount of coverage offered. The Affordable Care Act requires insurance companies to cover 10 essential healthcare services as well as increases transparency about rates and coverage between the insurance company and the insurance policyholder (Sebelius, 2013). Because of the Georgia Access Model, insurance companies take on more responsibilities in regards to driving enrollment[25]. As smaller insurance companies will not be able to keep up due to the lack of a centralized one-stop system (such as in healthcare.gov), larger insurance companies will have more market power to change prices. Insurance companies will be essential for providing smooth enrollment transitions, helping the state meet coverage enrollment goals, and providing transparency in a complicated system.
Rural Communities/People living in rural areas	Rural Georgians experience higher rates of uninsurance. With fewer provider options, it is more difficult to access care. As insurance companies will be required to conduct their own marketing, rural communities will feel the effect of these outreach efforts. However, due to the lack of an organized and centralized system, people may be steered towards plans that better benefit the insurer, not the purchaser. Georgia will not be able to match healthcare.gov's level of enrollment. With enrollment drops, hospitals will not be able to keep up with the strain that ensues on the system, leading to further hospital closures.

income eligibility limit to less than or equal to 138% of the FPL. By eliminating restrictive requirements and raising the income eligibility limit, more rural Georgians, who are more likely to be employed in sectors that are hard to document and would have more difficulty meeting the work requirements of Georgia's current program, will be able to obtain healthcare through Medicaid.

A traditional expansion of Medicaid, as opposed to the Georgia Pathways to Coverage program, will receive a higher federal matching rate to fund coverage, which would include 90% coverage of costs with an additional 5% for the first 2 years under the American Rescue Plan Act [18]. This expansion has already been enacted in 41 states, including many states that resemble the political makeup of Georgia [21]. With the government having already provided enhanced federal funding to so many states that have chosen to expand Medicaid coverage, Medicaid expansion in Georgia would be straightforward and legally permissible. Georgia's citizens, particularly those that are low-income and uninsured, would favor expansion because it gives them an opportunity to obtain much-needed healthcare coverage. Georgia and the CMS estimated that at least 447,000 people would be covered under traditional Medicaid expansion, which represents a 7-fold increase from the approximately 64,000 people that would be covered by Georgia Pathways [18].

While all proposed policy options have the potential to improve healthcare access in rural Georgia, they vary in feasibility. Table 1 summarizes that Medicaid expansion has the strongest impact but faces political barriers in Georgia. In contrast, strengthening outreach efforts as well as an AI-based oversight of insurance brokers score much higher across economic, legal, and political feasibility. Recent legislative initiatives and policy actions provide necessary context to evaluate the feasibility of the proposed policy solutions. The PeachCare Plus Act (SB 50) introduced by Senator David Lucas (D) and Sam Watson (R) aimed to provide subsidized healthcare coverage for low-income adults through a federal waiver, using federal funds to cover premiums and out-of-pocket costs. Concurrently, Senator Freddy Powell Sims (D) has highlighted the importance of health literacy and education amongst rural Georgians in her legislative work. The bipartisan support of these actions indicates that there is strong interest in a Marketplace-based expansion and leverage of federal funds as opposed to the traditional full Medicaid expansion. They portray the importance of policies that put consumer outreach and guidance first.

Final Recommendation

Georgia's Section 1332 Waiver and the Georgia Access Model were introduced with the goal of expanding coverage and increasing affordability. We strongly recommend a two-part corrective strategy to mitigate the harms imposed by the Georgia Access Model: (1) introducing a comprehensive outreach and education program for rural communities and (2) an Artificial Intelligence-powered oversight of insurance brokers and private enrollment platforms. Together, these address the primary shortcomings of the Georgia Access Model, which include fewer consumer protections and increased confusion about coverage options. A comprehensive outreach and education plan is vital to preventing coverage loss. High cost sharing and limited benefits are prohibitive to accessing care, but a coordinated statewide outreach strategy would be useful to improve health insurance literacy. However, education alone is insufficient. The AI-powered oversight provides a necessary enforcement that reduces the likelihood of substandard plans being selected. This plan would be successful in flagging non-ACA-compliant recommendations and prioritizing plans that maintain comprehensive coverage. The AI oversight is critical to restore consumer protection that was lost with Georgia's departure from the federal marketplace. An actionable plan to bolster health literacy and reduce the number of non-ACA-compliant plans among rural Georgians will significantly reduce the risk of vulnerable individuals purchasing the non-ACA compliant plans that are allowed by Georgia Access.

Though Medicaid expansion has the strongest potential impact on healthcare coverage, the premise of Waiver 1332 would imply a strong rejection of Medicaid expansion in Georgia's legislature. Given the low political feasibility of Medicaid expansion in Georgia's current political climate, we concluded that targeting outreach and education in addition to the implementation of AI-based oversight of brokers constitutes an effective solution. It is also important to be aware of the limitations of these proposals so that appropriate efforts can be taken to work towards these initiatives such that rural communities are not delayed getting access

to the best care they need. One such limitation would be that the proper development of an AI agent from data collection, model training, and regulatory testing may take years to unfold as this is typical of large scale public technological initiatives. However, this challenge can open doors for collaborations such as public-private partnerships and pilot programs in rural communities so that the technology is continuously implemented yet iteratively improved each year. There also exist bureaucratic or administrative barriers where party-line disagreements can lead to halts in progress or lack of funding to execute initiatives. This can be overcome by having advisory boards provide nonpartisan oversight of these proposals. Ultimately, the concerns regarding the economic and technical feasibility of these solutions would be outweighed by the outcome of success, which means increased insurance literacy, higher enrollment rates, and increased coverage among rural populations.

Rural Georgians do not have access to comprehensive coverage at an affordable price. Underneath the Georgia Access model, there remains the underlying issue of damaging consequences when individuals unknowingly purchase poor health insurance plans. Georgia Access aims to increase enrollment and lower premiums at the expense of comprehensive coverage, leaving rural Georgians with a lack of options in where to obtain healthcare. With the impending expiration of ACA premium tax credits, rural Georgians will be forced to choose between paying higher premiums for inadequate plans or forgoing insurance altogether, further straining rural healthcare systems that are already operating on thin margins. Improvement is a necessary action to prevent exacerbating rural health inequities. Ultimately, this two-part strategy that targets the flaws of Georgia Access would be the most realistic long-term strategy to ameliorate rural Georgia's limited access to health insurance.

Acknowledgements

We would like to express our gratitude towards Dr. Ashley Bradford for her guidance throughout our writing process. Her valuable feedback offered a new perspective on our research, and her comments helped to improve the clarity of our writing.

Access references for this paper using the QR code below.

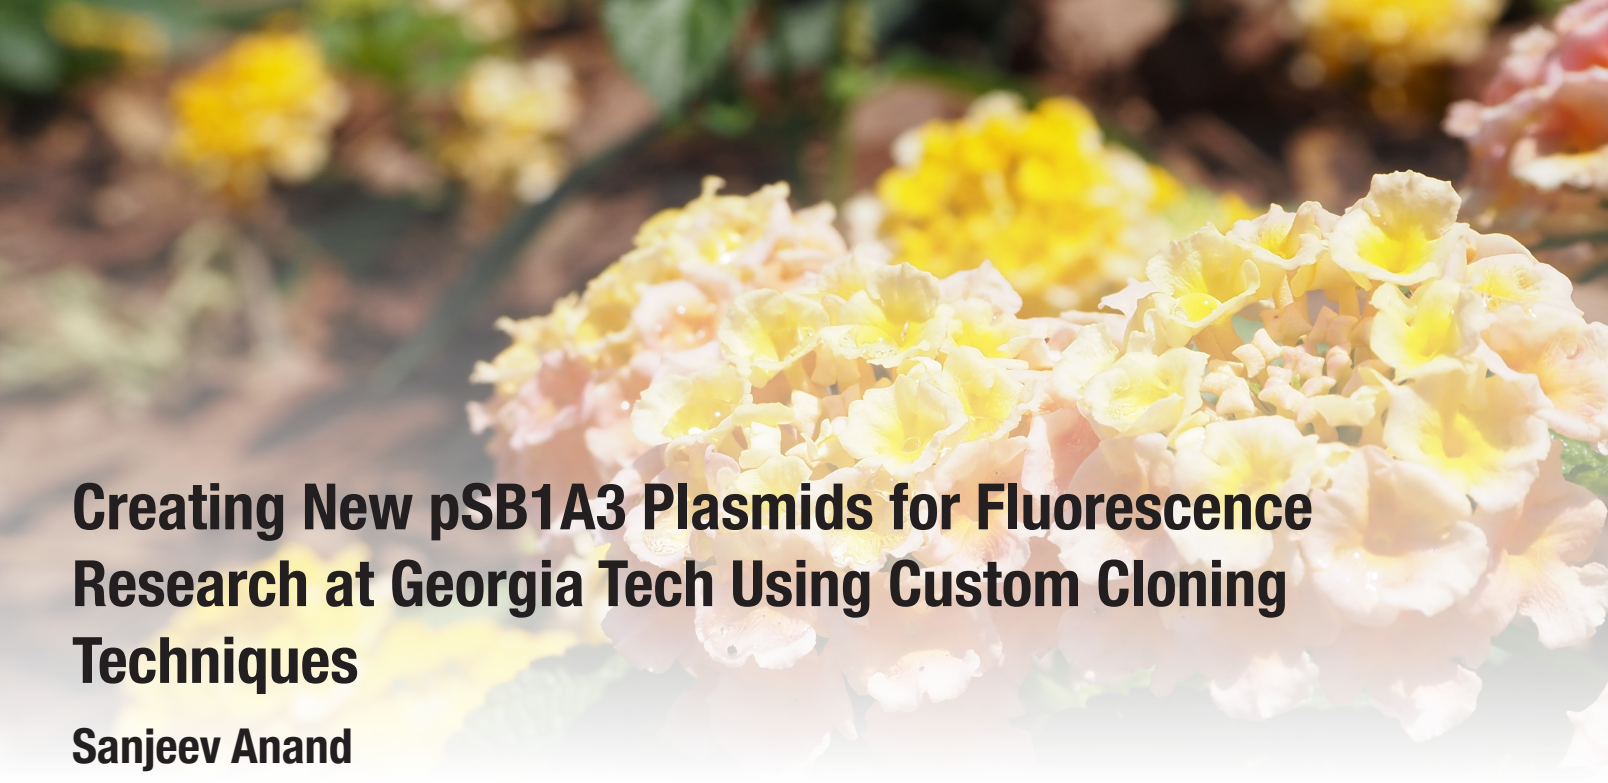

Creating New pSB1A3 Plasmids for Fluorescence Research at Georgia Tech Using Custom Cloning Techniques

Sanjeev Anand

Sanjeev Anand is a fourth-year biology major with minors in history and HMED currently on the pre-med and research track. He is currently a student researcher working under Dr. Anton Bryksin at the Molecular Evolutionary Core. He is currently interested in research involving fluorescent proteins and their use in the medical field. As someone who wants to work in the medical field, he is interested in detection methods such as those involving cardiomyopathic research. As a result, he hopes to be able to utilize the knowledge from his research in his own interests moving forward into medical school.

Abstract

Fluorescent proteins have been utilized for several decades here at Georgia Tech as a method for detecting fluorescence in industrial and biomedical research. This is accomplished by using a light at one wavelength to excite and later emit electrons at a different wavelength that is measured for the genes of interest. To be able to produce these wavelengths, a gene of interest is inserted into a plasmid of interest and then grown via custom cloning procedures, at which point they are confirmed using methods such as whole genome sequencing, fluorescence microscopy, and flow cytometry. These genes can be expressed through a variety of colors including red, blue, and green. As a result, my project has focused on developing five plasmids that can be utilized by the Molecular Evolutionary Core. To do so, I designed structures using a plasmid backbone, pSB1A3, which consists of an ampicillin resistance gene. I then designed an oligonucleotide synthesis of a promoter region consisting of two operators as well as a ribosomal-binding site which was later edited to be added onto primers of interest. After narrowing down which genes I would use, I performed Gibson Assembly on them alongside the backbone to discover primers that could be developed via PCR-based deletion and annealing procedures. From there, I created plasmid models that I plan to use when following through on my custom cloning procedures. Overall, this experiment serves as an opportunity for me to move forward with my research to physically design my plasmids of interest.

Introduction

Fluorescent proteins, or FPs, are proteins that conjugate to proteins of interest and express them using fluorescence. These proteins illuminate when an electron that was excited from a lower level to a higher level drops back down to the original level. This causes the release of a photon that is captured as a fluorophore. The fluorophores released are used to categorize the excitation and emission spectra of the FP, which serve as values of comparison between different models (Thermo Fisher Scientific, n.d.-c). Other comparisons include the fluorescence lifetime value (the amount of time after photon absorption that it takes to relax to the ground state) (FPBaseHelp, n.d.), half-life (how long a FP is functionable for) and maturation time (how long it takes for the protein to be created). They are used in a variety of settings, such as for detecting infarcted cells affected by cardiomyopathy and detection of sleeping sickness in patients (Zimmer, 2015).

FPs were originally discovered in Aequorea victoria jellyfish in 1961 by researchers at the University of Washington (Nikon, n.d.). This was done through the isolation of a bioluminescent protein called aequorin, which led to the discovery of a second protein called a green fluorescent protein (GFP). This protein illuminates with a green color when exposed to light. They are usually around 25,000 base pairs long (Yang et al., 1996). Over time, however, the gene was isolated and utilized as a probe. These probes included red fluorescent proteins (RFPs), a new family discovered in Discosoma anemone, as well as fluorescent probes, which are synthetically produced (Dietrich and Maiss, 2002). These RFPs helped create distinctions between different families of FPs as well as between proteins. Differences between FP families are usually found in the chromophore, or a group of three amino acid sequences responsible for the color of the protein that is expressed (Nikon, n.d.). However, modifications can be to other parts of the protein in order to change properties post-translationally (Shcherbakova and Verkhusha, 2014).

However, costs for fluorescent proteins have continued to rise. This is especially of utmost importance given the fact that materials for FP research can range from several hundreds of dollars to several thousands of dollars. As a result, creating them in house can be helpful. To be able to do so, one can look more into plasmid research as well as the processes required in order to derive these proteins. These proteins are coded by gene blocks that are inserted into a plasmid. Specific plasmids serve as the backbones for these proteins. These plasmids contain the usual structures, such as the origin of replication and antibiotic resistance genes (Rijal, 2024). However, considerations can be made in regards to the open reading frames, or areas where genes of interest can be inserted. This can be done through a method called Gibson assembly (Addgene, n.d.-b). This process consists of joining several DNA fragments together utilizing primers that were created for custom use through the utilization of databases. Two primers, usually greater than 20 base pairs, are required; one in the forward direction and one in the reverse direction (Yan, 2018). Following this process, PCR is used to amplify the DNA. The success of this process is attributed to the fact that DNA fragments can be combined regardless of size in a single tube. However, the type of plasmids used here also matter.

In order to select the plasmids for use, multiple considerations must be made. First, the number of base pairs in the open reading frames must be considered to allow enough space. Also, their copy number, or replicants they are able to create after cloning, must be considered (Cohrt, 2024). These can be divided into low (15-20) or high (500-700). While low copies are better able to compare to the wild type, high copies are able to create a large amount of compounds. Therefore, those that are high-copy are more likely to be considered.

A type of plasmid one I could consider is a version of the pUC19 plasmid, one of the most common plasmids utilized in research (Shetty et al., 2008). This research is conducted through Biobricks, a commercialized form of biological parts placed in the Registry of Standard Biological Parts. The base model usually consists of a base vector (BBa_I51020) with a high copy origin of replication and an ampicillin resistance marker. In order to change any of these components of the plasmid, one must utilize properties of both a type of cloning (Gateway cloning) and gene recombination. This can result in the creation of new vectors with alterations to its origin of replication or its antibiotic resistance genes. Several different combinations of these can be created, with a multitude of versions that can be developed for different processes. These versions are named

through the pSB system, a system based on plasmid synthetic biology, that follows the format pSB#X#, with the first number indicating the origin of replication, the letter indicating the antibiotic gene, and the second number indicating the version. Other part combinations can be added on using a dash and the addition of the part number.

To determine the version of plasmid, one can narrow down the origin of replication and the antibiotic gene. Although multiple versions can be considered, those that have a lack of mutations and a quick incubation time are considered to be optimal. For example, when choosing an antibiotic, ampicillin could be selected as it has a recovery period shorter than others despite the creation of satellite colonies (Stroik, n.d.). Similarly, when choosing an origin of replication, pMB1 could be selected due to the fact that it has a short size as well, thus allowing for a lack of mutations. Therefore, using the pSB system, one would be able to identify any plasmids with a code of pSB1AX (Shetty et al., 2008). This is due to the fact that 1 represents pMB1 in the naming convention, while A represents ampicillin. However, the last number remains variable, allowing one to choose the copy that they could use. One of these plasmid copies could be a plasmid such as pSB1A3 (Shetty and Knight, 2004). pSB1A3 is a low-copy plasmid with 2157 bp that can be cloned by restriction digest. This plasmid could thus be used for the backbone for any Gibson Assembly processes.

However, the inserts must also be considered if designing a research project. As stated earlier, these FP inserts arrive as g-blocks and can be inserted using Gibson Assembly. Therefore, in order to create models of interest, one must be able to discern different groups. It is always best to consider those within the same families as they usually contain the same properties. For example, some GFPs include CFPs (cyan fluorescent proteins), BFPs (blue fluorescent proteins), and YFPs (yellow fluorescent proteins) (Zimmer, 2015). In addition to this, some RFPs include Kaede, tdTomato, mCherry (mono Cherry), and mOrange (Dietrich and Maiss, 2002). Using algorithms such as Snapgene or Geneious, one would be able to identify key physical components to determine which inserts are able to meet certain criteria, such as fluorescent lifetime and spectra. This is critical given the fact that said properties will appear on any fluorescent imaging techniques such as plate reading and fluorescence microscopy when utilized in practice. Using the information about fluorescent gene proteins, one may be able to find some key examples to utilize in further research. These include eBFP, eCFP, meGFP, mCherry, mOrange, and tdTomato.

Therefore, one could use the information listed above to derive a model. Using the backbone listed above, pSB1A3, and the three plasmids from each of the green and red families, six plasmid models could be created to establish solid experimental models. This was something that was considered at the Molecular Evolutionary Core. In the past, one could have accomplished this with modules such as oligonucleotides and Sanger sequencing. However, new plasmids, especially one that provides immunofluorescence, can be especially helpful in the long run. Researchers across Georgia Tech would thus be able to utilize these plasmids moving forward, and such a stock could help our lab derive a new source of revenue moving forward. This would help set a precedent for future experiments and allow us to utilize additional open reading frames for methods of detection, thus helping to establish a new method of local testing for future clients. Although this project originally began as a student project, it could be extremely beneficial in order to allow better distribution of projects involving fluorescent proteins across campus. As a result, this novel concept can continue to be expanded to allow for additional pioneering research.

Question

Is there a way to add multiple fluorescent genes into a pSB1A3 plasmid backbone via Gibson Assembly and computer simulations?

Methods

First, identify the types of proteins that will be used. These must be selected from the two families (RFPs and GFPs). Then, select the proteins to be considered. Utilize a tree if possible to identify relatives and place them aside. Next, identify plasmids of interest. These must be compatible with the requirements listed in the introduction section, and may require multiple copies.

Following this, design an oligonucleotide that can serve as the starting and ending regions for any plasmids that are used. Most plasmids contain a terminator region already, so one is usually not required for this. However,

other regions to assist with the start of DNA replication and RNA transcription are required as these are usually manually inserted into the plasmid. Therefore, a collection of bases can be created using an oligonucleotide synthesis format. This involves designing a region with all necessary components and sending it out to the Molecular Evolutionary Core to be able to construct. To do so, the following format must be followed. A promoter, preferably a T5, must be attached to two downstream lac operators. From there, a 15-25 base region called a UTR spacer must be present before attaching to the Shine-Delgarno sequence of the oligonucleotide. Finally, a 6-9 nucleotide spacer must be present to separate the region from its attachment to the start codon. While the upstream region is attached to the Biobrick prefix, the downstream region is attached to the gene. These components can be attached using Gibson Assembly as well. This oligonucleotide must later be built into the primers of interest in order to facilitate an overlap region that allows for fewer Gibson Assembly products down the line.

Next, use pSB plasmids and design it using the naming convention established. To do so, determine the antibiotic and origin of replication of interest to identify a code to search up that would be useful in future experiments. From there, add the plasmid into Geneious alongside the sequence of the different FPs chosen. Moreover, make sure to run the Gibson Assembly, making sure to collect the files and data, including the primers of interest, to use in future experiments. This procedure also requires codon optimization to work in cells of interest, which requires using the Gibson Assembly software once again. Edit the primers as needed to contain the overlap regions listed above.

From there, make sure to purchase all plasmid blocks required for preparing both the plasmid backbone and the genes of interest. This procedure involves purchasing full plasmids from Addgene in order to excise portions using PCR methods. This is due to the fact that all of these follow the Biobrick base vector model (Shetty, et al., 2001). Using PCR techniques, attach primers to the ends of the sites. Anneal for both a backbone to be cut out as well as the genes inserted into the Biobrick vector sites. Use plasmid sequences either located in stock or from a company to create the genes of interest. Use an iGEM Kit Plate for the backbone (Biobrick BBa_K808000 well 14G is preferred).

Finally, create models that can be utilized for further experimentation. These will serve as checks to use once the custom cloning procedure is followed. These can be labeled and saved for further use, especially when used in combination with consensus sequencing that can be adapted once actual plasmids are created in the lab.

Results

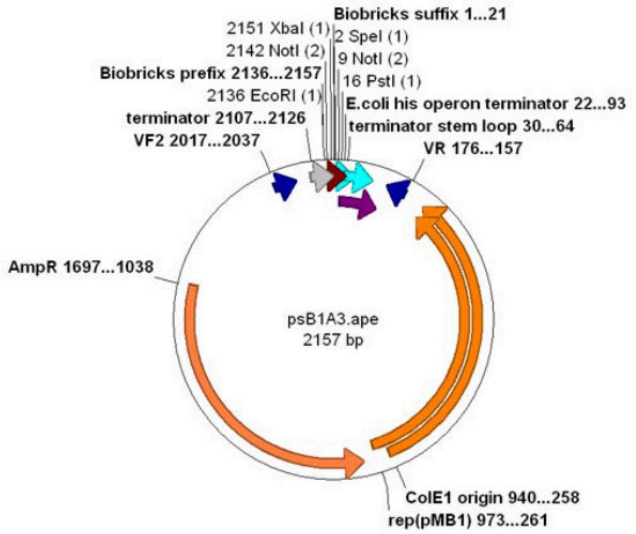

Figure 1. A depiction of the backbone derived from the pSB1A3 vector (Shetty et al., 2004). This includes various structures that must be considered,but could slightly vary in the region between the Biobrick prefix and suffix for the plasmid that was selected in house. However, this region is planned to be removed using molecular methods.

Protein	Number of Base Pairs
eBFP	720
eCFP	720
meGFP	720
tdTomato	1431
mCherry	711
mOrange	711

Table 1. A depiction of the genes of interest selected as well as the base pairs in the gene blocks produced. These include three from the green family and three from the red family.

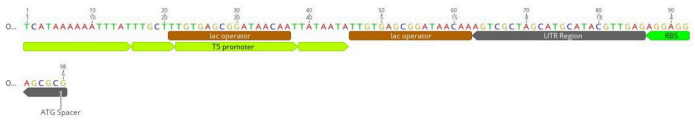

Figure 2. A depiction of the oligonucleotide synthesis sequence containing the promoters, operators, and RBS to be created (Dotmatics, n.d.). This synthesis is planned to be attached to the primers, as shown in the example below. This will allow for a decrease in the required number of Gibson Assembly pieces.

mOrange	pSB1A3	pSB1A3	Reverse	mOrange	GACTTTGT TATCCGCT CACAATAT TATAATTGT TATCCGCT CACAAAGC AAATAAAT TTTTTATG ACTCTAGA AGCGGCC GCGAATTC CAGAAATC ATC
	mOrange	pSB1A3	Forward	pSB1A3	TGGTATGG ACGAACTG TACAAATA GTACTAGT AGCGGCCG CTGCAGT CCGGCAAA AAA
	mOrange	mOrange	Reverse	pSB1A3	CAGCGGCC GCTACTAG TATCATTTG TACAGTTC GTCCATAC CACCGGTA GA
	pSB1A3	mOrange	Forward	mOrange	AACAAT TAATATTGT GAGCGGAT AACAAAG TCGCTAGC ATGCATAC GTTGAGAG GAGGAGC GCGATGGT TTCTAAAG GTGAAGA AAACAAC A TGGCG

Tables 2a.-b. A depiction of the primers used for an mOrange clone to attach to the gene of interest and backbone in order to anneal them moving forward. This synthesis is also planned to be created in our oligonucleotide synthesis department, thus allowing us to produce them in-house.

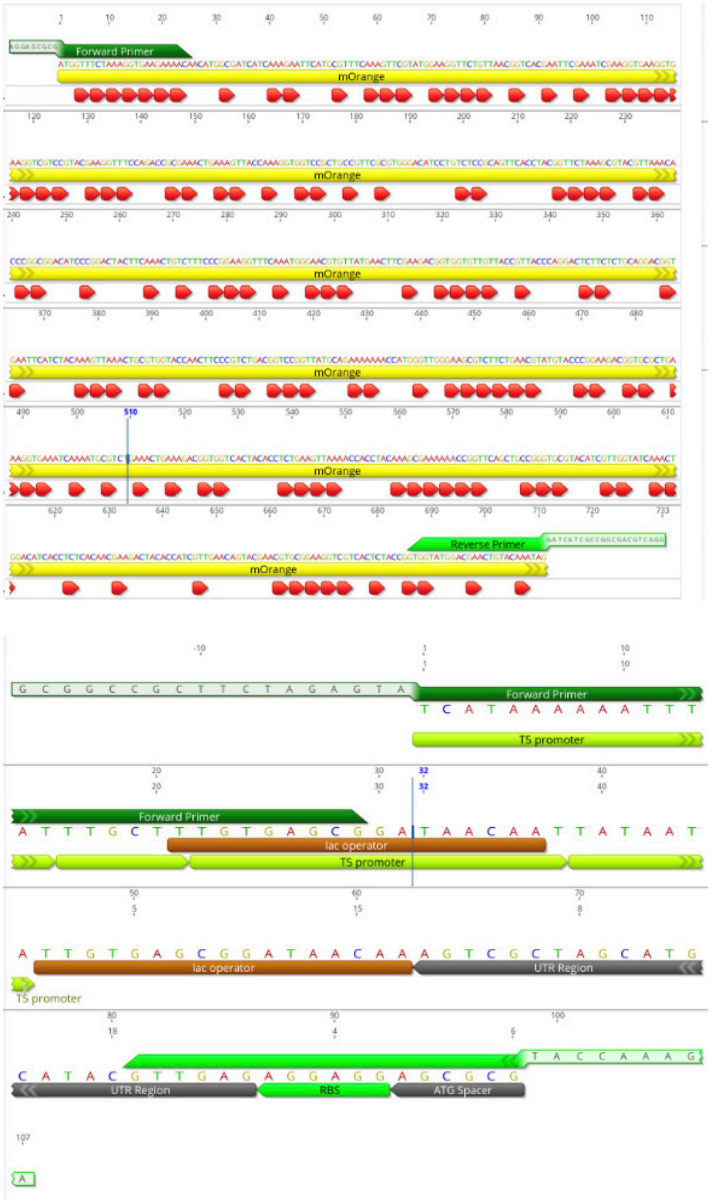

Figures 3a.-b. An example of the depiction of the PCR creations that are made using forward and reverse primers for mOrange (Dotmatics, n.d.). These were created from attaching the primers above to the backbone and gene of interest. While the genes of interest would need to be purchased, as stated above, the other parts could be created in house. As a result, most of the laboratory techniques that would follow could be conducted in the lab relatively quickly.

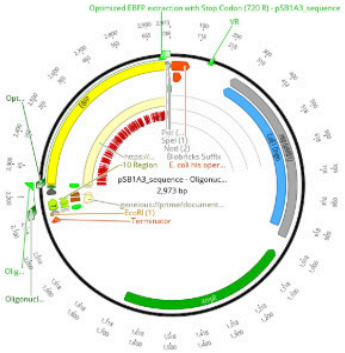

eBFP

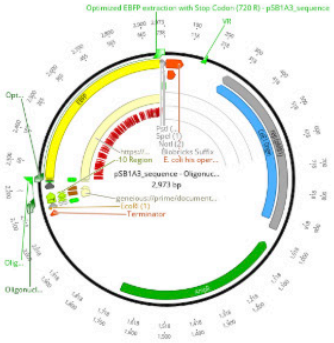

eCFP

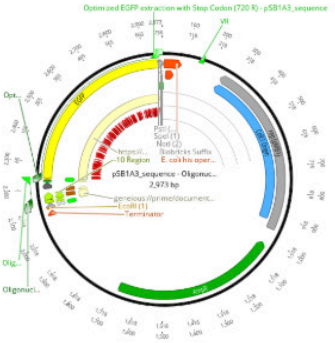

meGFP

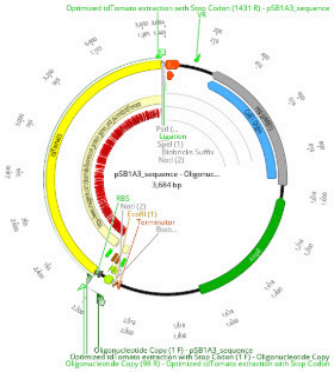

tdTomato

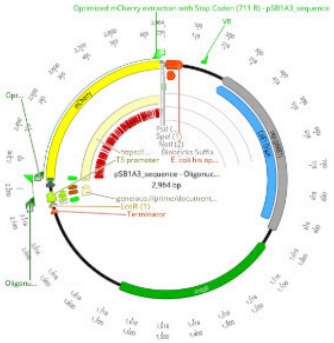

mCherry

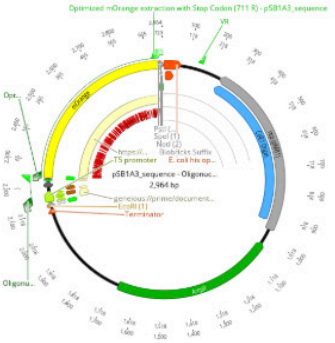

mOrange

Figures 4a.-f. A depiction of the plasmids that are to be produced by the MEC for future use. In the presentation, there is also a physical model of said plasmids added as well. 4a. eBFP, 4b. eCFP, 4c. meGFP, 4d. tdTomato, 4e. mCherry, and 4f. mOrange. The backbone for these plasmids could be purchased from sites like Addgene and then removed using further research techniques.

Discussion

I ended up creating six theoretical models of plasmids to use based on the pSB1A3 plasmid model. To do so, I first analyzed the plasmid of interest. Figure 1 depicts the pSB1A3 model with different components. I ended up finding out that these components are important to identify when creating our own plasmids of interest. This is pivotal given the fact that one technique I plan to use down the road is whole genome sequencing. From there, I moved onto identifying our genes of interest. I was interested in utilizing at least three red fluorescent proteins and three green fluorescent proteins as this would allow for a variety of proteins for our laboratory to use while also allowing for simple comparisons. Table 1 depicts these genes of interest used as well as their sizes. This will be important down the road given the fact that these genes will add to the sizes of the backbones when they are inserted using Gibson techniques.

From there, I was able to design the products necessary for Gibson assembly to take place. Figure 2 depicts the oligonucleotide synthesis that I ended up creating. This is due to the fact that a promoter region is not located in the Biobricks of choice. As a result, I would be able to help guide and control levels of transcription. Tables 2a.-b. utilizes this oligonucleotide synthesis and a theoretical gene of interest, mOrange, to simulate the process of designing primers. I chose to include one that enables us to shorten the size of our presentation, but also helps us to identify the components I would need when sending out such a design for production. I utilized these primers as well as the genes of interest in order to produce structures that would be created with further processes such as PCR amplification. Figures 3a.-b. serves as a depiction of the structures that would

be created.

Finally, Figures 4a.-f. show the final product that will be created once these plasmids are put together. They help to not only describe what they will look like down the road, but also serve as an important way to check their viability. This will be done using laboratory techniques such as whole genome sequencing to assure that mutations do not occur. Furthermore, they serve as a depiction of where genes are for other researchers to utilize when importing their own genes of interest. Overall, these plasmids serve as a great model to follow through on when designing them in the lab. Given that my research is still ongoing, I can utilize these models to order my genetic sequences of interest and prepare myself for further research down the road.

Conclusion

My research focused on the development of plasmids containing fluorescent genes for researchers at Georgia Tech. To do so, I focused on developing six plasmids, three containing genes of interest from the red family and three containing genes of interest from the green family. To test their efficacy, I utilized Gibson Assembly to simulate if they could be added to a backbone prepared from a pSB1A3 plasmid via a PCR method to attach primers. Overall, my theoretical models ended up discovering a way to create these six plasmids, allowing for such procedures to be created.

However, our work did not stop there. Although the theoretical concepts were prepared for further research, I still had to complete our physical work. Given that this research is an ongoing procedure, I will be conducting wet lab processes moving forward. However, identifying these backbones, primers, genes of interest, and oligonucleotide synthesis only serve as the first step. While I prepared for my research to continue, I produced a methodology. This was pivotal to understanding what materials I would need to gather moving forward while also allowing me to follow through on replicating my work for further trials.

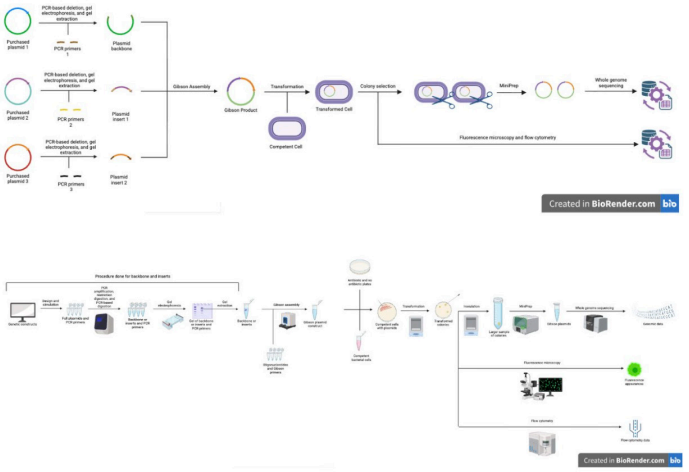

Figures 5a-b. A depiction of the entire methodology to be followed when developing fluorescent plasmids containing new genes of interest (Biorender, n.d.). This includes techniques that I will follow moving forward in my research, allowing me to visualize my goals.

Figures 5a-b. depicts the procedural and theoretical techniques that I would need to follow for my plasmids to be expressed and tested for. To do so, this requires first preparing all of the stocks that I would need, including pouring plates and making cells competent. Afterwards, I would need to identify the backbone of the plasmid I plan to utilize, which would need to be amplified using custom cloning procedures. At the same time, I would also need to order plasmids with my inserts of choice alongside the primers that I am interested in using. These primers that I design would need to be annealed to my fragments of interest via PCR before being measured via gel electrophoresis and extracted using gel extraction. From there, a custom cloning model would need to be repeated to produce additional copies of the new plasmids before testing them using techniques such as whole genome sequencing, fluorescence microscopy, and plate reading. These results would then be analyzed and published, allowing researchers to utilize them.

Overall, I expect that the plasmids that will be created will be useful for Georgia Tech researchers. This procedure is considered to be long, but the payoffs will be immense. I anticipate being able to produce a new source of plasmids researchers at Georgia Tech will be able to use moving forward. By creating a new source of plasmids, we at the Molecular Evolutionary Core hope to help in the development of cutting-edge fluo-

rescent protein research in laboratories all around us.

Acknowledgements

First, I would like to thank my research PI, Dr. Anton Bryksin, and my research mentor and second reader, Naima Djeddar, for all their help in this project. They were my guides throughout this project and helped me with all the requirements that I needed to publish my thesis. Second, I would like to thank all the other researchers who I worked with at the Molecular Evolutionary Core. This includes, but are not limited to, the Sanger Sequencing laboratory, with Adam Fallah, and the Custom Cloning module, with Ashtyn Cauffman, Cameron MacKay, and Adarsh Nallapa. Third, I would like to thank my former research mentor, Nicole Diaz, who assisted me with the project when I first started in fall 2024. I also would like to thank my LMC professor, Dr. Courtney Hoffman, for her help in helping me design my project proposal. Fourth, I would like to thank my research advisor and BIOS 4699 professor, Dr. Chrissy Spencer, for her assistance in helping me to meet the guidelines of a four-semester research project. Finally, I would like to thank my mother, Priya Srinivasan, who consistently read my papers when I had immediate doubts, and helped me in whatever way she could to clarify what I could do to improve based on her background in research knowledge.

Access references for this paper using the QR code below.

Build a Career WITH PURPOSE

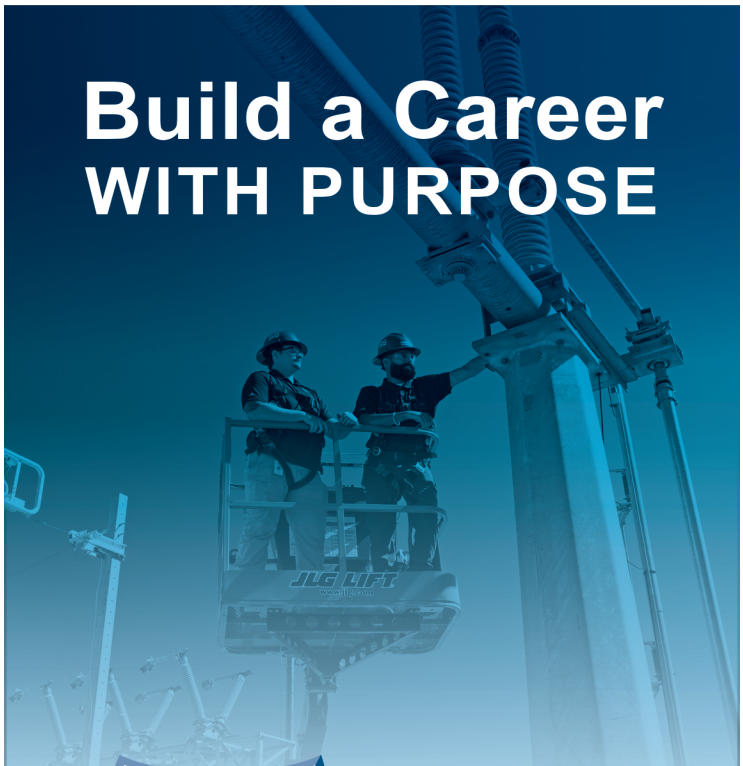

Southern States

Powers the Future™

Build Your Future with Southern States

Work alongside experienced professionals, gain hands-on experience, and grow through mentorship and meaningful projects shaping the power industry. Scan the code to apply today!

SouthernStatesLLC.com



ENGINEERED FOR SUCCESS

At Georgia Institute of Technology, success is engineered from the ground up. That's why Georgia Tech trusts the Miura LX Modular Boiler System to provide reliable dry steam for its cutting-edge research and development labs. Like many prestigious facilities around the world, Georgia Tech relies on Miura for safe, efficient, and reliable steam solutions. When precision and performance are non-negotiable...

MIURA BOILERS ARE THE ONLY CHOICE.

miuraboiler.com

MIURA



Developing and Validating Quiet Quitting Intention: A Theory-Based Framework of Intentional Effort Regulation

Anish Cheekati

Anish Cheekati is an undergraduate researcher in organizational psychology at Georgia Tech working under the supervision of Dr. Chris Wiese, in collaboration with Yichen Tang. His research focuses on employee motivation, effort regulation, and contemporary workplace behavior. He plans to continue research in organizational psychology and explore applications of motivational theory to modern work design.

Abstract

Quiet quitting has attracted widespread attention (Fuller & Kerr, 2022; Zenger & Folkman, 2022), yet scholarly definitions and valid measures remain limited. Popular accounts often frame quiet quitting as disengagement or boundary-setting, conflating it with burnout, withdrawal, or low performance. Drawing on the Theory of Planned Behavior, we conceptualize Quiet Quitting Intention (QQI) as the deliberate motivation to limit discretionary effort while continuing to meet formal job requirements. Across three studies, we developed and validated the QQI-15, a psychometrically sound measure capturing both Extrinsic Quiet Quitting Intention (EQQI) and Intrinsic Quiet Quitting Intention (IQQI). Study 1 used qualitative interviews, theory-driven item generation, and expert review to establish content validity. Study 2 refined the scale through exploratory and confirmatory factor analyses in large, independent samples, yielding a five-factor structure that loads onto two higher-order dimensions. Study 3 established construct, discriminant, and incremental validity within a broader nomological network, demonstrating that QQI is distinct from burnout, disengagement, and counterproductive work behavior while predicting meaningful work outcomes beyond existing quiet quitting measures. By centering intention rather than behavior, this research clarifies the psychological foundations of quiet quitting and provides a theoretically grounded tool for studying intentional effort regulation in contemporary work.

Existing academic efforts to measure quiet quitting have primarily operationalized the construct behaviorally or attitudinally without clearly distinguishing motivational intention from enacted effort restriction (Anand et al., 2024; Atalay & Dağıstan, 2024). In contrast, the present framework centers intention as the focal construct and differentiates between extrinsic and intrinsic referent standards guiding effort regulation. This distinction allows for clearer differentiation between deliberate boundary-setting and affective withdrawal, representing a conceptual refinement over existing measures.

Introduction

The concept of *quiet quitting* has gained substantial traction in both public discourse and academic inquiry, particularly in the wake of widespread organizational disruption following the COVID-19 pandemic. Popular accounts typically describe quiet quitting as employees doing the required amount, declining to go above and beyond formal job duties while remaining employed. Although this framing has intuitive appeal, it has also generated conceptual ambiguity. Quiet quitting has been variously described as disengagement, early-stage withdrawal, burnout, or resistance to hustle culture, often without clear differentiation among these constructs (Anand et al., 2024; Atalay & Dağıstan, 2024).

A central limitation of existing work is the tendency to define quiet quitting primarily in behavioral terms. Observable reductions in discretionary effort—such as refusing optional tasks or limiting availability—are treated as direct indicators of the phenomenon. However, similar behaviors may arise from fundamentally different psychological processes. Employees may reduce extra-role effort due to emotional exhaustion, lack of motivation, perceived unfairness, or deliberate boundary-setting rooted in personal values. Without attending to the underlying intention behind effort regulation, behavioral definitions risk collapsing distinct motivational states into a single construct.
The present manuscript addresses this limitation by shifting the focus from quiet quitting behavior to Quiet Quitting Intention (QQI). QQI is defined as a deliberate cognitive orientation toward restricting discretionary effort while continuing to meet formal role expectations. This distinction aligns with longstanding research in organizational psychology demonstrating that intentions are conceptually and empirically distinct from behaviors and often serve as their most proximal antecedents (Ajzen, 1991, 2002). By centering intention, QQI captures the motivational decision-making process that precedes or accompanies observable effort regulation.

Building on this foundation, we propose a multidimensional framework that differentiates between two forms of quiet quitting intention based on their guiding standards. Extrinsic Quiet Quitting Intention (EQQI) reflects effort regulation anchored in external referents, such as job descriptions, reciprocity norms, or perceptions of fairness. Intrinsic Quiet Quitting Intention (IQQI) reflects effort regulation anchored in internal referents, such as personal values, well-being, or identity-based boundaries. This referent-based distinction provides a more precise account of why employees may choose to limit discretionary effort and clarifies how quiet quitting differs from adjacent constructs. Emerging scholarship has started to explore the concept of quiet quitting using attitudinal or behavioral measures (e.g., Anand et al., 2024; Atalay & Dağıstan, 2024). However, existing definitions operationalize quiet quitting as withdrawal or disengagement behaviors without placing an emphasis on the underlying motivational intention. The present adds to the existing literature through emphasizing intention as the core construct.

The goals of the present research are threefold. First, we develop a theory-grounded conceptualization of QQI that distinguishes intention from behavior. Second, we construct and validate a psychometrically sound measure of QQI, capturing both EQQI and IQQI. Third, we situate QQI within a nomological network of related constructs to establish its discriminant validity. In doing so, this work contributes to a more rigorous and nuanced understanding of quiet quitting as a form of intentional effort regulation.

Theoretical Framework

Quiet Quitting as an Intention-Based Construct

The distinction between intention and behavior is central to theories of motivated action. Intentions reflect deliberate decisions shaped by beliefs, evaluations, and perceived constraints, whereas behaviors represent enacted outcomes that may or may not fully reflect those intentions. In workplace contexts, employees may harbor intentions to reduce effort that are not immediately observable due to situational constraints, fear of consequences, or normative pressures. Conversely, employees may exhibit reduced effort behavior without intending to disengage, such as when constrained by workload or resource limitations.

Conceptualizing quiet quitting as an intention rather than a behavior resolves several inconsistencies in prior definitions. It allows for the identification of effort regulation before it manifests behaviorally, differentiates voluntary restraint from involuntary withdrawal, and aligns quiet quitting with established models of decision-making. Accordingly, QQI is defined here as

a deliberate intention to limit discretionary effort while continuing to fulfill formal role requirements.

Theory of Planned Behavior

The Theory of Planned Behavior (TPB; Ajzen, 1991, 2002) provides a structured framework for understanding how such intentions form. TPB posits that intentions are shaped by three components: attitudes toward the behavior, subjective norms, and perceived behavioral control. When applied to quiet quitting, this framework suggests that employees form intentions to restrict effort when they evaluate restricting effort as beneficial or justified, perceive it as socially acceptable or normative, and believe they can do so without undue negative consequences.

TPB is particularly well suited for modeling QQI because it explicitly separates motivational antecedents from behavioral enactment. This separation allows for a more precise examination of effort regulation as a cognitive decision rather than a passive outcome of exhaustion or disengagement.

Extrinsic and Intrinsic Quiet Quitting Intention

While TPB explains how intentions form, it does not specify the standards used to evaluate their behavior. To address this gap, we distinguish between extrinsic and intrinsic referents that guide effort regulation.

Extrinsic Quiet Quitting Intention (EQQI) reflects an intention to regulate effort based on externally defined standards. Employees high in EQQI may believe that limiting effort is fair, justified, or necessary given their job description, compensation, or treatment by the organization. This form of intention aligns closely with social exchange processes, wherein effort is calibrated in response to perceived reciprocity or contract fulfillment (Blau, 1964).

Intrinsic Quiet Quitting Intention (IQQI) reflects an intention to regulate effort based on internally defined standards. Employees high in IQQI may prioritize personal well-being, autonomy, or life balance and intentionally restrict discretionary effort to maintain these values. This form of intention aligns with self-determination perspectives emphasizing autonomy and self-regulation.

Although EQQI and IQQI are conceptually distinct, they are not mutually exclusive. Employees may experience both simultaneously or shift between them over time. Distinguishing these forms of intention is essential for understanding the diverse motivational pathways that underlie quiet quitting.

Methods

Overview of Studies

Scale development and validation followed best-practice guidelines for construct development (Hinkin, 1998; Worthington & Whittaker, 2006; Holmbeck & Devine, 2009) and occurred across three studies. Study 1 focused on item generation and content validation. Study 2 refined the item pool and established factorial validity through exploratory and confirmatory factor analyses. Study 3 examined the nomological network of QQI and tested incremental validity relative to existing quiet quitting measures.

Study	Sample	Purpose	Method	Key Outcome
Study 1	N = 28 employed adults	Item generation & construct validation	Semi-structured interviews + thematic analysis	Thematic structure informing initial item pool
Study 2	N = 741 (EFA n = 370; CFA n = 371)	Factor structure & scale refinement	EFA + CFA	Final QQI-15 validated
Study 3	N = 639 employed adults	Construct & incremental validity	Correlational and regression analyses within nomological network	Discriminant & incremental validity supported

Table 1. Overview of Study Objectives, Samples, and Analytic Approaches

Study 1: Item Development

Study 1 employed qualitative interviews, theory-driven construct mapping, and expert review to generate an initial item pool. Interviews with self-identified quiet quitters were conducted to capture authentic motivational experiences underlying effort regulation. Participants were recruited online through professional networks and were screened to confirm they were employed full-time. Interviews were semi-structured and were conducted for 30-60 minutes, including prompts such as: "How do you define going above and beyond at work?", "What motivates you to restrict discretionary effort?", and "How do you determine what level of effort is appropriate?" All interviews were recorded and transcribed. An interactive thematic analysis approach was used for coding with initial codes generated in-

dependently by trained coders and then reviewed for consistency and theoretical alignment. Thematic analysis yielded six theoretically derived subdimensions, using procedures consistent with Braun and Clarke (2006), corresponding to attitudes, subjective norms, and perceived behavioral control for both extrinsic and intrinsic forms of quiet quitting intention. A combined human- and AI-assisted item generation process produced an initial pool of 60 items, which were evaluated for clarity, relevance, and content validity through expert review and cognitive interviews.

Study 2: Scale Refinement and Factor Structure

A preliminary pilot study (N = 93) assessed item clarity and response variability, resulting in minor wording revisions and the removal of two weak items. For the primary validation phase, data were collected from 741 employed adults, randomly split into independent samples for exploratory factor analysis (EFA; n = 370) and confirmatory factor analysis (CFA; n = 371), following recommendations for exploratory and confirmatory factor analysis (Bandalos & Finney, 2018). EFA supported a six-factor structure aligned with theorized subdimensions. Subsequent CFA rounds compared competing models and progressively refined the scale.

Across three CFA rounds, items demonstrating low loadings or redundancy were removed. Self-Regulatory Control failed to demonstrate adequate convergent and discriminant validity and was dropped. The final model retained 15 items across five first-order factors (Effort–Reward Imbalance, Perceived Work Norm, Boundary Control, Self-Fulfillment Orientation, and Personal Norm), was loaded onto two higher-order dimensions: EQQI and IQQI. The final higher-order model demonstrated excellent fit (CFI = .970, RMSEA = .061). The higher-order EQQI and IQQI latent factors were moderately correlated (r = .58), which would imply that the dimensions are related but not mutually exclusive constructs.

Study 3: Nomological Network and Incremental Validity

Study 3 examined construct validity using a new sample of 639 employed adults recruited via Prolific and assigned to a planned-missing design, consistent with established standards for construct validation (Campbell & Fiske, 1959; Cronbach & Meehl, 1955).

The finalized 15-item Quiet Quitting Intention scale (QQI-15; see Appendix A for full scale items) demonstrated strong convergent validity with theoretically related constructs (e.g., job satisfaction, engagement, perceived organizational support) and discriminant validity from counterproductive work behaviors and demographics. Incremental validity analyses showed that EQQI and IQQI explained additional variance in key outcomes beyond existing quiet quitting scales, supporting the unique contribution of an intention-based conceptualization.

Results

Factorial Validity (Study 2)

In Study 2, exploratory and confirmatory factor analyses supported a multidimensional structure corresponding to theorized QQI subdimensions. Confirmatory analyses demonstrated that multifactor models substantially outperformed one- and two-factor alternatives. The final retained model consisted of five first-order factors loading onto two correlated higher-order constructs, EQQI and IQQI. This model showed strong global fit (CFI = .970, RMSEA = .061) and clear empirical separation between extrinsic and intrinsic intention pathways.

Reliability

Across Studies 2 and 3, the five QQI-15 subscales demonstrated acceptable to strong internal consistency. Cronbach's alpha coefficients for the subdimensions ranged from .66 to .84 across samples. Specifically, internal consistency estimates for the subscales fell between .663 and .827, exceeding conventional thresholds for research use (Cortina, 1993). These reliability estimates were consistent across studies and analytic approaches.

Construct and Incremental Validity (Study 3)

In Study 3, nomological analyses supported the validity of QQI as a distinct motivational construct. QQI was strongly associated with exchange-related variables (e.g., perceived organizational support, psychological contract breach), moderately associated with burnout and engagement, and negligibly related to counterproductive work behavior. Turnover intention was more strongly related than originally hypothesized, suggesting it may represent a downstream outcome of quiet quitting intention rather than a discriminant construct. Importantly, QQI demonstrated incremental validity

over existing quiet quitting measures, particularly in predicting job satisfaction, engagement, organizational citizenship behavior, and perceived organizational support.

Discussion

The present research advances understanding of quiet quitting by reframing it as an intention-based construct rather than a purely behavioral outcome. Across three studies, we developed and validated the QQI-15 and demonstrated that quiet quitting intentions are multidimensional, comprising distinct extrinsic and intrinsic pathways that reflect different standards for regulating discretionary effort. This reconceptualization clarifies longstanding ambiguity in the literature and helps explain why prior work has yielded inconsistent findings when quiet quitting has been operationalized solely through observed behaviors (Anand et al., 2024; Atalay & Dağıstan, 2024).

Interpreting the Higher-Order Structure of QQI

A central contribution of this work is the identification of a hierarchical structure in which five first-order facets load onto two higher-order dimensions—EQQI and IQQI. This structure indicates that employees may arrive at similar patterns of effort regulation through qualitatively different motivational routes. EQQI reflects effort calibration grounded in external standards, such as perceived reciprocity, role expectations, and fairness judgments. In contrast, IQQI reflects effort calibration grounded in internal standards, including personal values, identity-consistent boundaries, and well-being priorities. Importantly, these pathways are not simply different labels for disengagement; they represent deliberate, reasoned decisions about how much discretionary effort to invest at work.

The presence of Boundary Control as a bridging facet further underscores this interpretation. Boundary management appears to function as a mechanism through which both extrinsic and intrinsic standards translate into intention. For some employees, boundaries are enacted in response to perceived overreach or inequity; for others, boundaries reflect proactive self-regulation aligned with personal goals. Recognizing this shared mechanism helps reconcile debates about whether quiet quitting is reactive resistance or proactive boundary-setting—it can be either, depending on the underlying intention.

Distinguishing QQI from Related Constructs

By centering intention, QQI is conceptually and empirically distinct from burnout, disengagement, and counterproductive work behavior. Whereas burnout reflects emotional exhaustion and disengagement reflects reduced psychological investment, QQI captures a deliberate decision to withhold discretionary effort while maintaining role performance. The moderate associations observed between QQI and burnout suggest that emotional strain may co-occur with quiet quitting intentions for some employees, but the weaker associations relative to engagement and exchange-related constructs indicate that QQI is not reducible to affective depletion.

Similarly, the negligible relationships between QQI and counterproductive work behavior highlight an important distinction between withholding extra-role effort and engaging in harmful or deviant actions. Quiet quitting intention reflects restraint, not sabotage. This distinction has practical significance, as conflating quiet quitting with misconduct risks misdiagnosing the underlying motivational state and responding with inappropriate interventions.

Turnover intention emerged as more strongly related to QQI than originally hypothesized, suggesting that it may represent a downstream outcome rather than a competing construct. This pattern is consistent with models of withdrawal in which employees progressively recalibrate effort before exiting the organization. From this perspective, quiet quitting intention may function as an early indicator of disengagement trajectories, offering organizations a window for intervention prior to turnover.

Implications for Theory

The QQI framework contributes to organizational theory by integrating the Theory of Planned Behavior with motivation and social exchange perspectives to explain intentional effort regulation. Prior research has often treated effort withdrawal as a passive response to strain or dissatisfaction. The present findings demonstrate that effort regulation can be intentional, structured, and value-driven, emphasizing the role of cognitive evaluation and perceived control in shaping work behavior.

This intention-based approach also helps reconcile

disparate findings in the emerging quiet quitting literature. Behavioral measures alone cannot differentiate between employees who reduce effort due to exhaustion, those responding to perceived injustice, and those proactively setting boundaries to preserve well-being (Maslach et al., 2001; Demerouti et al., 2001). By distinguishing EQQI and IQQI, the present framework provides a clearer lens for interpreting observed behaviors and for developing cumulative theory.

Practical Implications

The distinction between extrinsic and intrinsic quiet quitting intentions has important implications for organizational practice. Because EQQI and IQQI represent related but distinct motivational pathways, employees may vary in the degree to which each dimension characterizes their effort regulation. Employees higher in EQQI, reflecting effort restriction grounded in perceived unfairness or exchange imbalance, may respond more favorably to interventions aimed at restoring fairness, clarifying role expectations, or improving exchange relationships. In contrast, higher IQQI reflects intentional boundary-setting anchored in personal standards, which may not be responsive to traditional engagement initiatives. Importantly, these dimensions are not mutually exclusive; individuals may exhibit elevated levels of both, depending on contextual and motivational factors.

Misclassifying IQQI-driven effort regulation as disengagement may lead organizations to overcorrect, inadvertently violating employees' boundaries or eroding trust. Conversely, failing to address EQQI may allow perceptions of inequity to persist, increasing the likelihood of withdrawal and eventual turnover. Diagnosing the intentional basis of effort regulation is therefore critical for designing effective, targeted responses.

Limitations and Future Directions

Several limitations warrant consideration. All studies relied on self-report data, which may introduce common method variance. Although the intention-based focus mitigates some concerns by capturing cognitive orientations rather than behaviors, future research should incorporate multi-source and longitudinal designs. Longitudinal studies would be particularly valuable for examining transitions between IQQI, EQQI, and downstream outcomes such as disengagement and turnover.

Future research should also explore contextual antecedents of QQI, including organizational policies, leadership behaviors, and cultural norms surrounding availability and overwork. Additionally, examining supervisor perceptions of quiet quitting intentions—and potential misattributions—may shed light on performance evaluation processes and manager–employee dynamics.

Theoretical Implications

This work contributes to organizational theory by integrating TPB with motivation and exchange perspectives to explain quiet quitting. It extends prior research by demonstrating that effort regulation can be intentional, structured, and value-driven rather than reactive or dysfunctional.

Practical Implications

Understanding whether employees' effort regulation is driven by extrinsic or intrinsic intentions has practical implications for organizational intervention. Efforts to address quiet quitting may require different responses depending on whether employees are reacting to perceived unfairness or seeking to protect personal well-being.

Limitations and Future Directions

The present research relies on self-report data and cross-sectional designs, which limit causal inference. Future research should examine longitudinal dynamics of QQI and explore how organizational practices influence shifts between EQQI and IQQI over time.

Conclusion

Quiet quitting is best understood not as disengagement or burnout, but as a form of intentional effort regulation. By centering intention and distinguishing between extrinsic and intrinsic standards, the Quiet Quitting Intention framework provides a clearer theoretical and empirical foundation for studying this phenomenon. The validated QQI-15 scale offers researchers and practitioners a tool for distinguishing deliberate boundary-setting from emotional withdrawal, advancing both theory and practice in the study of contemporary work behavior.

Acknowledgements

The author would like to thank Yichen Tang for valuable collaboration and intellectual contributions throughout the development and validation of the Quiet Quitting Intention framework. The author also expresses gratitude to Dr. Chris Wiese for mentorship and guidance during the research process. Appreciation is extended to all study participants who contributed their time to this project.

Appendix A

Quiet Quitting Intention Scale (QQI-15)

The QQI-15 consists of two subscales: Extrinsic Quiet Quitting Intention (EQQI) and Intrinsic Quiet Quitting Intention (IQQI). Items assess respondents' intentional decisions to regulate discretionary effort while continuing to meet formal role requirements. Items are rated on a Likert-type scale indicating agreement with each statement. Higher scores reflect stronger quiet quitting intentions.

Access references for this paper using the QR code

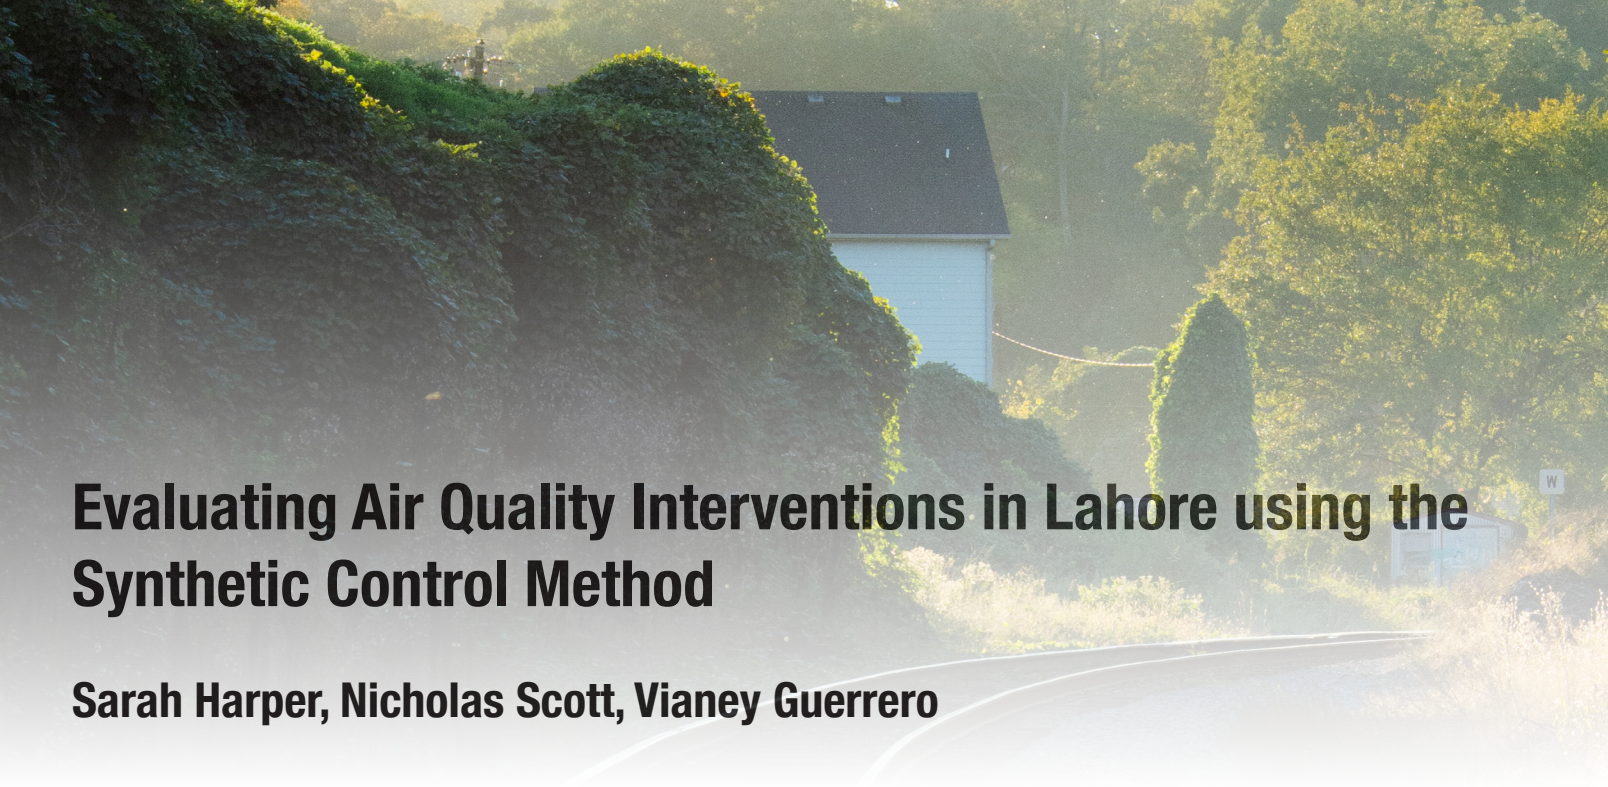

Evaluating Air Quality Interventions in Lahore using the Synthetic Control Method

Sarah Harper, Nicholas Scott, Vianey Guerrero

Vianey Guerrero is a third-year undergraduate student Industrial Engineering major with a Public Policy minor. Her academic interests focus on applying data driven problem solving to improve systems while integrating policy considerations to create efficient and socially responsible solutions.

Nicholas Scott is a second-year undergraduate Computer Science major. His academic interest focuses on using data science to improve quality of life in the fields of environmental sustainability, energy, and sports analytics.

Sarah Harper is a fourth-year undergraduate International Affairs major with an Applications of Artificial Intelligence and Machine Learning minor. Her academic interests focus on bringing together data science and policy to examine human rights issues across Asia.

Abstract

Air quality is a critical crisis within Pakistan, specifically Lahore, and this is due to mass crop burning. This study examines the effect of crop burning bans within different agricultural periods from 2017-2020, specifically focusing on 2019 when Pakistan implemented its crop ban. We theorize that AQI will be moderate during and increase after the ban due to low compliance, enforcement, and pressing agricultural deadlines. However, within this study, we find evidence counter to the intervention suggesting the ban did not reduce noncompliance and was associated with increased AQI levels. Thus, after the month-long ban was lifted, there was a rebound effect due to the quick turnaround time to clear the fields with burning for planting season.

Motivation and Background

Lahore's crisis-level air quality has frequently been amongst the worst in the world. AQI (Air Quality Index) in the Punjab region, where Lahore lies, is incredibly high, causing public respiratory health issues, school closures, and extreme smog (Waqas et al 2024; "School Closures in Pakistan" 2024). While multiple factors contribute to this pollution, this study focuses on crop residue burning in agricultural areas around Lahore due to the politically contentious debates surrounding it. Crop burning is detrimental to the AQI, because it releases particulates such as PM2.5 and PM10 into the air, which in turn affects the AQI. PM 2.5 is especially hazardous as the particulates are small enough to penetrate deeply into the lungs and enter the bloodstream (IQAir Staff Writers 2024). These effects are why understanding the effect of crop burning on Lahore's AQI is crucial. Crop burning has become a largely debated environmental policy in Pakistan, as it is a widespread issue with limited solutions. Farmers argue that it is a cheap and fast way for farmers to clear their fields before the new planting season. Whereas environmentalists, residents, and the Pakistani government argue that the thousands of fires and health issues require alternative solutions. This debate has increased various policy implications, such as burn bans, implementing fines, and mechanical alternatives for crop stubble removal, like the Super Seeder (Malik 2025). This project will specifically examine burn bans because, although a month-long ban will be implemented in Pakistan in 2019, there appears to be no significant effect on mitigating crop burning (Shaikh 2022). Understanding why these bans are ineffective through examining their effect on AQI can help develop effective policy interventions. Therefore, this study aims to answer the research question: What is the effect of crop burning bans on air quality in Lahore, Pakistan? This study contributes to the understanding of environmental policy issues in Lahore by examining impact of a crop burning ban on air quality. The findings highlight policy and behavioral challenges that showcase broader implications of crop burning regulations for similar regions. This project will employ a synthetic model to examine the effects of crop burning on AQI and will use NASA satellite data to track fires in agricultural areas near Lahore to examine what policy interventions would suit the situation best. The remainder of the paper is organized as follows. The next section reviews existing literature regarding the history of crop burning in Lahore and similar cities. The following section explains our hypotheses based on the literature review. The next three sections detail our data collection plan, sources used, and the result of the data collected. We conclude with an action plan that details policy implications for improving air quality outcomes based on the results discovered.

Literature Review

To address the question of how crop burning bans affect AQI in Lahore, it is crucial to examine the existing literature in three areas: first, the general effect of crop burning near Lahore, and second, the crop burning bans' successes/failures in other countries, both before and after crop burning bans. This section analyzes other nations because most literature within Pakistan overlooks how the harvesting cycle interacts with crop burning bans and their effects on AQI. This literature review aims to create a literary foundation upon which this paper's analytical framework will be constructed.

In addressing the general effect of crop burning near Lahore, scholars examining this area argue that crop burning is aligned with the planting season as it is a way to clear fields, as well as that it is a huge contributor to Pakistan's air quality crisis. Azhar et al. (2019) examine the effect of crop burning on AQI before any ban implementations from 2000 to 2014 by examining emissions from wheat, rice, maize, and sugarcane once burned. The study found that rice straw was the largest contributor to NH_3 and NO_x and that sugarcane was the largest contributor to black carbon (BC), two of the crops that are grown and burned closest to Lahore, Pakistan (Azhar et al 2019). The study also found that burning was concentrated in the peak harvest months, depending on the type of crop. Azhar et al. argue that crop burning is a significant and growing contributor to its air quality crisis, along with that provinces like Punjab are primary areas for burning due to their agricultural infrastructure. Scholars like Anwar (2025) agree with these conclusions, as their study examined crop residue burning and air pollution in Punjab also found that there were patterns around planting seasons. Their research on crop burning data from 2020 to 2024 found that peak burning was from October to February and correlated with farmers clearing the land in

preparation for planting season (Anwar et al 2025).

However, while these scholars provide valuable information on the correlation between planting schedules, AQI, and crop burning, they do not evaluate the effect of crop burning bans on AQI. Scholars like Khanna and Cao, while examining nations other than Pakistan, conclude that within China and India, bans are ineffective due to poor enforcement and compliance. Khanna examines India's crop burning ban implemented in 2015 and argues that while there was an initial effectiveness, once the public saw that the ban was not being enforced, they continued to burn. This study found that the enforcement level correlated with whether fires continued to burn or not. Areas where the fines were low and rarely enforced had more fires. This study demonstrates that bans are not effective, however, it is that with poor enforcement, they are likely to be disobeyed (Khanna et al 2025). Cao's study about straw burning bans in China from 2002-2019 also found that burning moves to months of weaker enforcement instead of being eliminated (Cao et al 2023). The study also found that while nighttime burns were easier to constrain, it was harder to enforce daytime burns and more rural agricultural areas due to cost limitations. Both Cao and Khanna argue that limited changes in burning behavior are a result of poor enforcement.

Scholars such as Erbaugh et al. and Lan et al. argue that burning returns to normal once a burn ban is lifted, due to poor enforcement, quick turnaround times for planting season, and no other cheap/fast alternative to burning. Erbaugh et al. surveyed farmers in India and found that many reverted back to burning again after the bans due to the high costs of no-burn machinery ($2,100-$3,500 plus higher diesel costs). They also found that farmers burned due to the short window of time to manage crop residues before planting, and that they found enforcement tactics to be more of an annoyance rather than effective (Erbaugh et al 2024). Lan et al. support these findings as their 2022 study found that the 2009 Punjab Preservation of Subsoil Water Act—meant to conserve water during summer months–delayed rice planting dates to correlate with monsoon rains. However, this then caused only 10-20 days between rice harvest and wheat sowing, pressuring farmers to burn to clear fields (Lan et al 2022). Both studies argue that due to the quick turnaround time, farmers return to burning regardless of whether there is a ban or not. Alternatively, scholars like Robbins, examining Indian natural environments, argue that resistance patterns can occur depending on the perception of the state and political interventions and whether they are legitimate or not (Robbins 1998). This is supported by Costa and Khan examining environmentalist ideology and finding that ideology affects responses to environmental investments and that conservatives are less likely to participate in or support environmental improvements to social policy. This further supports the argument that defiance and a negative response can be an outcome of environmental nudges like crop burning (Costa and Kahn 2013).

Hypothesis and Theory

While the existing literature provides valuable explanations on the general effect of crop burning and bans within other nations, there is a severe lack of literature addressing the crop ban in Pakistan and its effect on AQI. Based on the existing literature discussed above, this study has three theories.

The first hypothesis addresses what happens during a crop burning ban. Due to non-compliance and low enforcement during crop burning bans documented in Khanna and Cao, from 2023 and 2025, crop burning bans will not affect AQI. This hypothesis is justified by the fact that when enforcement and compliance are low, the ban functions more as a nominal practice rather than effective prevention. Because of this, burning continues as normal, perhaps at a slightly decreased rate depending on the enforcement level. This then results in moderate AQI levels; however, while it may be better than unrestricted burning, it would be worse than if burning had stopped completely.

Therefore, we hypothesize that,

H1: *Weak enforcement and low compliance will have minimal effect on crop burning behavior from October to November, resulting in similar baseline AQI levels in Lahore.*

We are also interested in examining behavior after the ban is lifted. One possibility from arguments in Erbaugh et al. (2024), Lan et al. (2022), Azhar et al.(2019), and Anwar (2025), is that due to non-compliance and pressing planting seasons, crop

burning increases after a ban. This hypothesis is justified by the fact that the planting cycle begins in April, and that although burning may have subsided in October and November, this does not discount the need for farmers to clear the fields. This results in an accumulation of non-burned agriculture and causes farmers to mass-burn after the ban is lifted, resulting in a large increase in AQI.

Therefore, after the ban is lifted, we might expect a rebound effect. This leads to our second hypothesis that,

H2: *Planting schedules and low compliance will cause farmers to resume burning when the November ban ends to prepare fields for the April planting cycle, resulting in a significant increase in Lahore's AQI.*

An alternative possibility is that crop burning could actually increase after the ban is lifted due to defiant behavior. Defiance behavior has been documented in multiple settings, such as Robbins (1998), Costa and Khan (2013), which supports our theory that crop burning will increase due to defiance behaviors. This creates situations where crop burning is an act of enforcement/political resistance, inferring that the amount of burning across the period would have happened without the ban, as defiance behavior becomes more popular. Therefore, instead of reducing burning or maintaining non-ban levels, the ban worsens the issue altogether. When farmers see the poor enforcement and minimal repercussions, as suggested in our first hypothesis, the perceived risk of crop burning decreases. This then increases the social norm of deliberately participating in non-compliance, as it is more normalized within agricultural communities

This leads to our third hypothesis that,

H3: *Defiance behavior will cause crop burning to increase from October to April, resulting in a significant increase in AQI levels in Lahore.*

Our preferred theories are hypotheses 1 and 2, which predict that during the ban, there will be no change in crop burning behaviors due to non-compliance and poor enforcement, and that after the ban, crop burning will increase due to a limited time to prepare for the upcoming crop season and low compliance.

Alternatively, other explanations for changes in AQI include things like agricultural economics that reflect increased crop production versus decreased production. Additionally, crop burning pollution in neighboring countries like India could cross over and affect the Pakistani AQI. Other causes include brick kilns, factories, or vehicle emissions–all causes that fluctuate with different seasons. However, due to time and data constraints, this study will not test these alternative hypotheses. Through these hypotheses and alternative explanations, this study aims to understand the effects of burn bans in Lahore to provide more effective policy solutions.

Methodology

This study evaluated the effectiveness of the one-month crop burning ban in Lahore in 2019 using the Synthetic Control Method (SCM). The intervention period is defined using the official start date of the ban and verified with NASA's Fire Information for Resource Management System (FIRMS) (2025) satellite fire directions to ensure alignment with the seasonal burning cycle (Map 1 & 2).

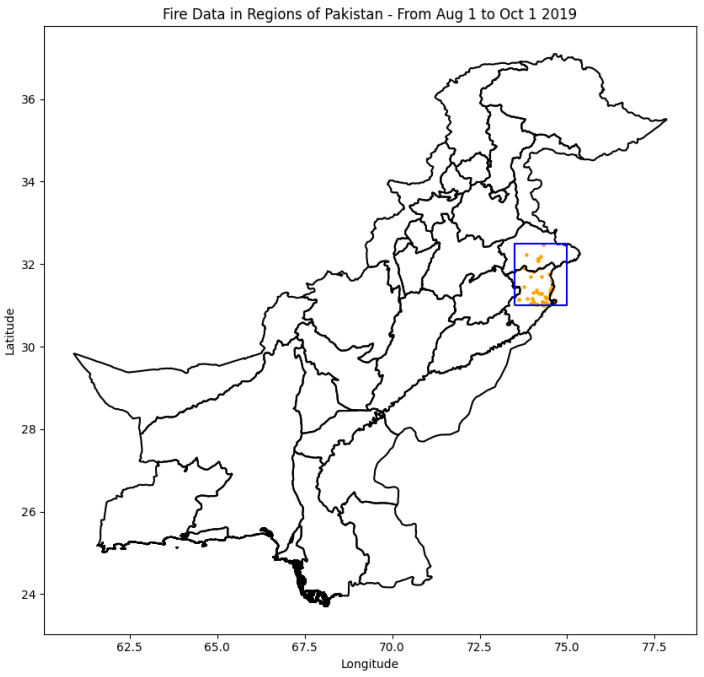

Map 1 cont.

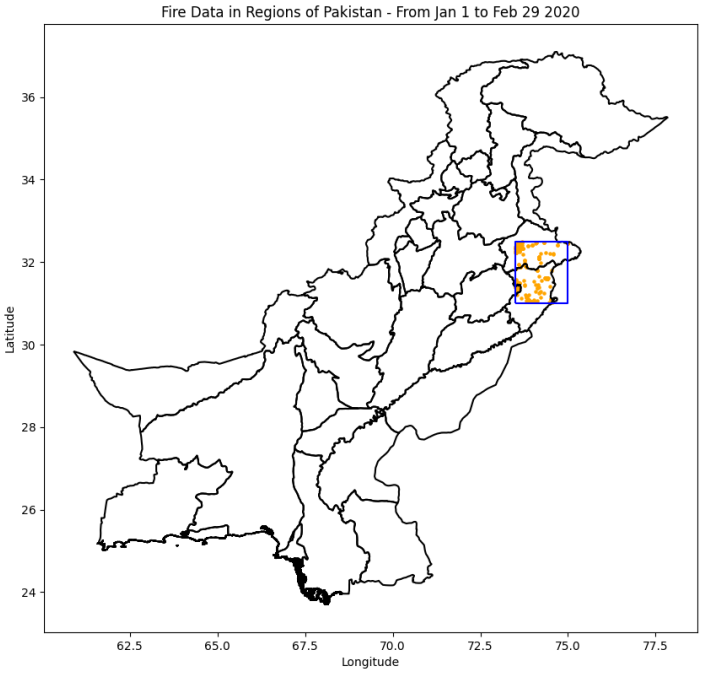

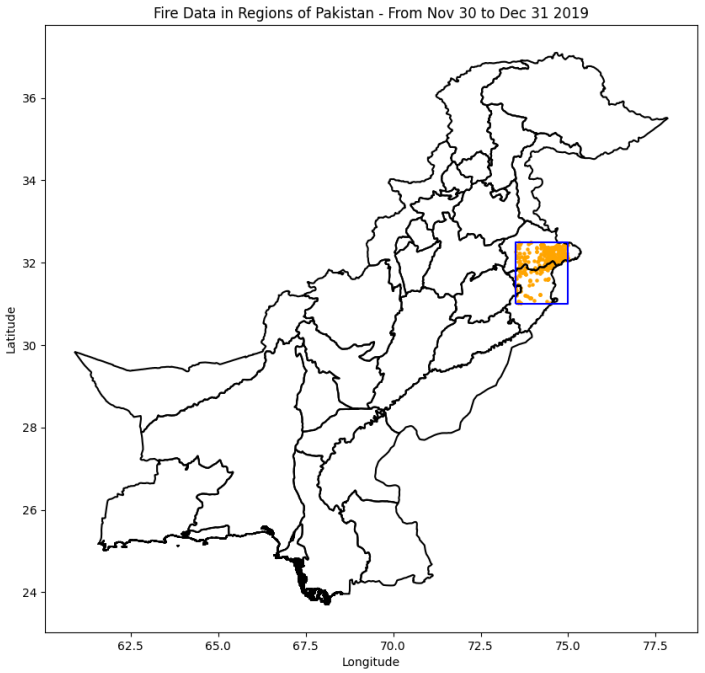

Map 1: Fire incident locations in the region surrounding Lahore during non–crop-burning months. The left map shows fire detections from August 1 to October 1, 2019, and the right map shows detections from January 1 to February 29, 2020. The highlighted area outlines the Lahore vicinity, indicating that recorded fires in both periods are not associated with the typical crop-residue burning season.

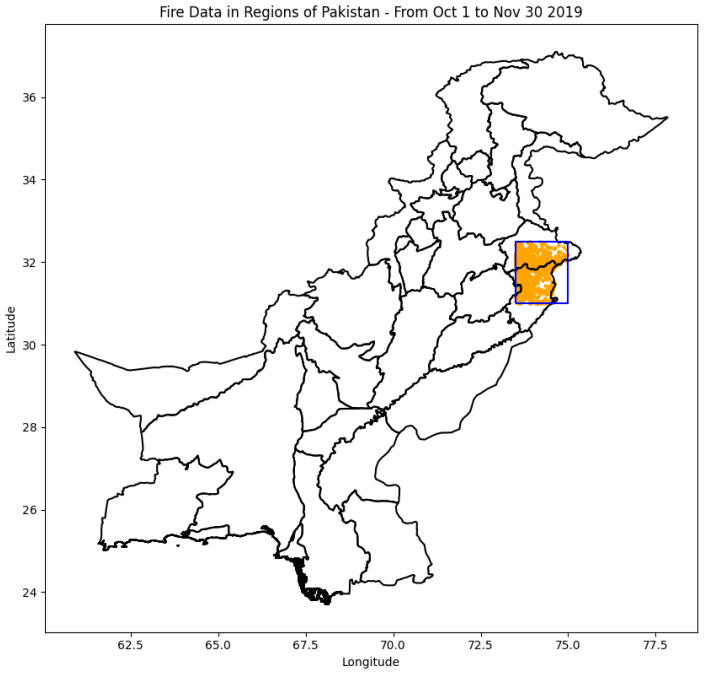

Map 2: Fire incident locations in the region surrounding Lahore during the late crop-burning season. The left map shows fire detections from October 1 to November 30, 2019, and the right map shows detections from November 30 to December 31, 2019. The highlighted area outlines the Lahore vicinity, indicating that fire activity is concentrated within this region during the peak crop-residue burning period.

Lahore is treated as the intervention unit. To estimate the air quality trajectory that Lahore would have experienced in the absence of the ban, a synthetic control was constructed from a weighted combination of donor cities selected for their similarity in pre-intervention air quality, temperature, humidity, and precipitation characteristics. We refer to this as Synthetic Lahore.

The predictor balance table (Table 1) shows that the weighted average of key predictors between Lahore and its synthetic counterpart is closely matched, indicating a strong pre-treatment fit. The donor weights (Appendix Item 1) show that Peshawar (56.1%) and Dhaka (0.149%) are the cities contributing the largest weights in the donor pool. This composition supports the assumption that these cities serve as an appropriate counterfactual for Lahore during the pre-ban period.

Predictor Variables	Weights	Treated_Real	Synthetic_Est	Difference	Pct_Difference
Temperature	0.78	72.00	66.920	5.079999	7.1%
Humidity	0.22	62.75	57.875	4.874997	7.8%
Precipitation	0.00	8.25	33.375	-25.125002	-304.5%

Table 1: Predictor Balance Table

A credible counterfactual requires close pretreatment matching between the treated unit and its synthetic counterpart. Table 1 summarizes the balance between Lahore and synthetic Lahore on key predictors. The synthetic control achieves a close match on temperature and a moderate match on humidity, though precipitation shows as a greater imbalance, which contributes to a weight of zero. Predictor weights (Table 1) show that temperature and humidity contributed most heavily to the model's fit.

The synthetic control method compared Lahore's actual air quality during the ban month with the synthetic control's predicted air quality under a scenario of no intervention. To assess whether the estimated treatment effect could happen by chance a placebo in space test was conducted by reassigning the treatment to both donor cities. The gaps are shown in the Placebo-in-space test model (Appendix Item 3). The purpose of this test is to evaluate whether Lahore's post-intervention diverges more clearly from its synthetic control than those of placebo units from theirs, which strengthens confidence in the counterfactual design.

To reduce bias, we used multiple pre-treatment years to strengthen counterfactual construction and reduce noise. As well as placebo testing (Appendix Item 3) to assess whether estimated effects exceed those obtained by chance. This identifying strategy relies on assumptions. Seasonal patterns in meteorological variables are accounted for by the model. Donor cities represent the air quality determinants relevant to Lahore.

Data Sources and Measures

The data used in this analysis come from two primary sources. Fire activity is measured using NASA's FIRMS (2025) which provides satellite-derived active fire detections essential for capturing the intensity and timing of crop burning events during the October - November fire seasons from 2017 to 2020. Meteorological variables are obtained from the WeatherDirectory (2025), which compiles weather data observations like temperature, humidity, and precipitation. Air quality measurements are drawn from a curated Dropbox repository containing datasets from U.S. Department of State (2024) consulates and embassies. Because the Synthetic Control Method requires complete, uninterrupted time series, all datasets undergo an extensive cleaning process that identifies missing values, removes data outside the analysis window, and fills gaps using averaging techniques. After cleaning, AQI values are log-transformed in R to stabilize variance and enable comparability across years. Summary statistics for the treatment and the donor cities further illustrate the structure of the data. The treatment sample shows relatively stable AQI and meteorological conditions, compared to the donor pool, which exhibits greater variability across cities (Appendix Items 4 and 5).

Results

The addition of new donor cities strengthens the trends seen in the plots and is supported by the percent improvement, which is 100% (Appendix Item 2). This table also shows that the SCM is stable with no measurable bias, given the average estimated bias being zero. The p-value is 0.36, which indicates that the intervention ban had no statistically significant impact on air quality. This supports hypothesis 1 that the ban had little impact on burning behavior. The visualizations in the SCM (Plot 1) and placebo-in-

space (Appendix Item 3) provide similar evidence. Focusing on the SCM, the pre-trend years show an overlap, while during the intervention period the normalized AQI is steady with a slight increase. With an increase in AQI post ban seen more in Treated than Synthetic. This provides validation for hypothesis 2 that an increase in AQI will occur post ban.

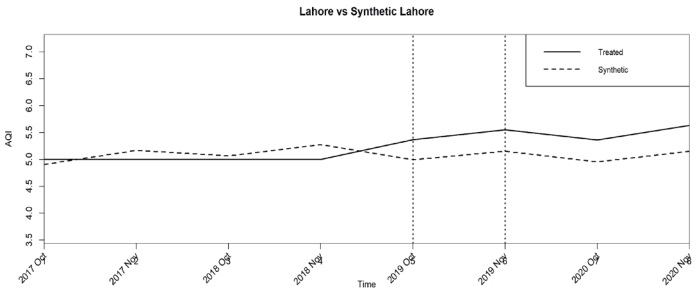

Plot 1: Synthetic Control Method (SCM)

In the event-study plot (Plot 2), the error bars during and post treatment period are near zero. This indicates a weak positive association between the treatment and normalized AQI levels, suggesting a slight increase in AQI levels following the treatments. This also supports hypothesis 2.

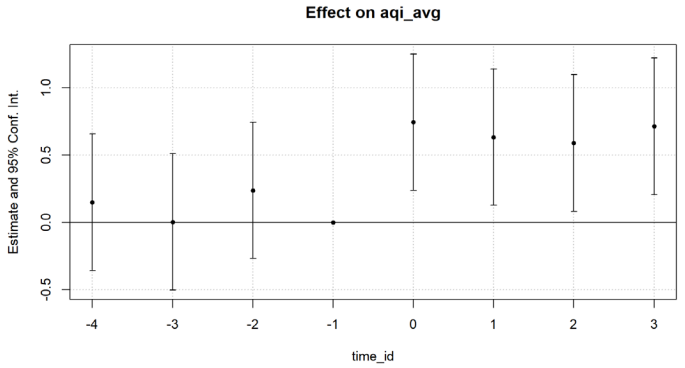

Plot 2: Event Study Plot

Thus, from our results, the intervention didn't impact crop burning behavior, supporting primarily hypothesis 1. Nonetheless, there were noticeably higher AQI levels for the treatment city, which increased during and after the ban, providing some support for hypothesis 2.

Action Plan

These results show that short-term restrictions are insufficient to reduce Lahore's AQI during the crop burning season. Because crop residue burning is staggered across farms and takes place across several weeks, policies must combine longer coverage, a stable legal framework, and credible alternatives for farmers. The following section examines three key policy implications that could improve air quality outcomes: extending burn bans to cover the full crop residue season, establishing a permanent legal framework with predictable enforcement, and promoting the adoption of sustainable alternatives such as subsidized super seeders and shared machinery schemes.

1. Extend the Seasonal Burn Ban to Cover the Full Harvest Season

An issue is that farmers delayed burning until after it became clear that enforcement was low (Shaikh 2022). Extending the ban to cover the entire burn season, which is typically 6 to 10 weeks, would reduce the probability of post-ban burning surges. Integrating technology and enforcement ensures that the ban is both credible and effective, reducing seasonal spikes in pollution.

An extended ban should be paired with real-time monitoring, including satellite detections and rapid response enforcement teams. Clearer timelines improve predictability for farmers, while continuous coverage prevents shifts in burning behavior that raise AQI levels

2. Establish a Permanent Crop Residue Burning Law with Standardized, Graduated Penalties

Burning prohibitions in Punjab are issued under Section 144, which

"Grants the provincial and district authorities unchecked power to issue directions to persons to either abstain from certain conduct or to direct persons to engage in certain actions for specific periods of time." (Hussain 2024)

A permanent statute would allow policymakers to integrate crop residue burning restrictions with agricultural planning and long-term environmental targets. To increase deterrence without imposing burdens, the law should adopt a graduated penalty structure:

- **First offense**: Warning plus mandatory training on approved crop residue management techniques.
- **Second offense**: Monetary fines.
- **Third offense onwards**: Higher fines, possible equipment confiscation or temporary land use restrictions.

Standardizing penalties reduces arbitrary enforcement,

enhances procedural fairness, and supports compliance.

3. Expand Access to Viable Alternatives Through Subsidies and Training

Regulatory bans on crop residue burning alone cannot succeed unless farmers have access to practical, affordable alternatives. In Punjab, the government has begun distributing subsidized super seeders, which can sow wheat directly into rice stubble, reducing the need for burning. However, the program faces significant challenges. Currently, there are only about 1,300 super seeders in the entire country (Rahul 2024) , creating a scarcity that drives up prices and limits access for most farmers. Even with a subsidy, the cost of the machine remains $1,800 USD (PKR 508,832), which is a substantial investment for small farms. Effective use of these machines requires a certain level of technical knowledge, and many farmers lack experience, making adoption slow and uneven.

To address these issues, the government should encourage larger, wealthier farms to help offset the cost of super seeders for smaller family farms, for example, cooperative ownership models or machinery sharing. In addition, farmer training programs and demonstration plots remain essential to teach proper operation and highlight the economic benefits of residue management. Combining financial support through cooperative schemes with knowledge transfer not only builds capacity but also could create a scalable, community-driven transition away from crop burning, improving air quality, and maintaining agricultural productivity.

Broader Considerations for Implementation

These recommendations require strengthened agency coordination between environmental, agricultural, and local government institutions as well as targeted investment in monitoring infrastructure and rural outreach. Partnerships with village-level vigilance committees can further support accountability and early detection of violations. Finally, improved data collection, such as fire-grained spatial records of burn events, would enable adaptive management and help refine policy design over time.

Conclusion

Overall, the analysis demonstrates that the short-term crop burning ban had minimal impact on farmers' burning behavior and only a slight temporary increase in air quality indicators, which highlights the limited effectiveness of brief interventions. This study was subjected to various limitations. The short duration of the ban made it difficult to ensure the air quality results were not affected by any outside factors. Another limitation was the lack of literature on the specific month-long crop burning. Future studies should examine more years of seasonal data and use more detailed monitoring networks to distinguish the bans' impact from external factors. Additionally, exploring farmers' adoption of alternative residue management practices and assessing the socio-economic feasibility of long-term bans will be crucial for developing effective, sustainable strategies to mitigate smog in Lahore.

Acknowledgements

We sincerely thank Dr. Omar I. Asensio, Director of the Data Science and Policy Lab for his guidance. We also acknowledge Reagan Lengefeld for her support.

Appendix

Appendix Item 1.

	Weights	Unit Name	Unit Id
1	0.032	Kathmandu	1
2	0.149	Dhaka	2
3	0.034	Chennai	3
4	0.031	Hyderabad	4
5	0.050	Mumbai	5
6	0.070	New Delhi	6
7	0.561	Peshawar	7
8	0.074	Islamabad	8

Weights of Donor Cities

Appendix Item 2.

```
Average ATT Estimate (p Value for Joint Null):  0.481   ( 0.36 )
L2 Imbalance: 0.000
Percent improvement from uniform weights: 100%

Avg Estimated Bias: 0.000

Inference type: Conformal inference

 Time Estimate 95% CI Lower Bound 95% CI Upper Bound p Value
    5    0.455             -0.521             1.431   0.198
    6    0.566             -0.410             1.542   0.213
    7    0.353             -0.623             1.328   0.217
    8    0.549             -0.427             1.524   0.198
```

Augsynth Average Treatment Effects

Appendix Item 3.

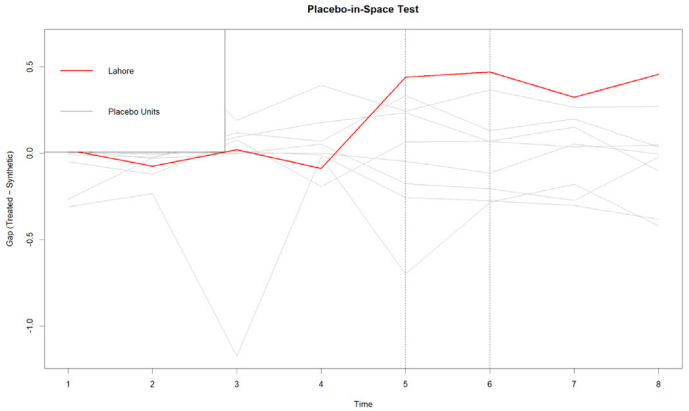

Placebo-in-Space Test

Appendix Item 4.

```
    city              city_id       year         month             aqi_avg          temp
Length:8           Min.   :0   Min.   :17.00   Length:8           Min.   :5.000   Min.   :71.00
Class :character   1st Qu.:0   1st Qu.:17.75   Class :character   1st Qu.:5.000   1st Qu.:71.38
Mode  :character   Median :0   Median :18.50   Mode  :character   Median :5.181   Median :72.00
                   Mean   :0   Mean   :18.50                      Mean   :5.239   Mean   :72.00
                   3rd Qu.:0   3rd Qu.:19.25                      3rd Qu.:5.414   3rd Qu.:72.62
                   Max.   :0   Max.   :20.00                      Max.   :5.635   Max.   :73.00
    humid            precip          time_id
Min.   :56.00   Min.   : 0.00   Min.   :1.00
1st Qu.:58.00   1st Qu.: 4.75   1st Qu.:2.75
Median :61.50   Median : 9.00   Median :4.50
Mean   :61.75   Mean   : 8.75   Mean   :4.50
3rd Qu.:63.50   3rd Qu.:13.00   3rd Qu.:6.25
Max.   :72.00   Max.   :19.00   Max.   :8.00
```

Treatment Summary Statistics

Appendix Item 5.

```
    city              city_id         year         month             aqi_avg
Length:64          Min.   :1.00   Min.   :17.00   Length:64          Min.   :3.803
Class :character   1st Qu.:2.75   1st Qu.:17.75   Class :character   1st Qu.:4.883
Mode  :character   Median :4.50   Median :18.50   Mode  :character   Median :5.000
                   Mean   :4.50   Mean   :18.50                      Mean   :4.990
                   3rd Qu.:6.25   3rd Qu.:19.25                      3rd Qu.:5.152
                   Max.   :8.00   Max.   :20.00                      Max.   :5.734
     temp            humid           precip          time_id
Min.   :56.84   Min.   :46.00   Min.   :  0.00   Min.   :1.00
1st Qu.:68.36   1st Qu.:65.75   1st Qu.:  9.00   1st Qu.:2.75
Median :75.83   Median :74.00   Median : 33.50   Median :4.50
Mean   :74.19   Mean   :71.70   Mean   : 73.11   Mean   :4.50
3rd Qu.:81.36   3rd Qu.:80.00   3rd Qu.:112.50   3rd Qu.:6.25
Max.   :85.28   Max.   :86.00   Max.   :366.00   Max.   :8.00
```

Counterfactual Summary Statistics

Access references for this paper using the QR code below.



Evidence and Neurological Mechanisms Behind Sugar Addiction

Deandra Smith

Deandra is a Neuroscience major with a minor in Health and Medical Sciences at the Georgia Institute of Technology. Her research interests include neurodegenerative disease, neuropharmacology, scientific communication and outreach, and data-driven approaches in neurological and physiological sciences. Following completion of her undergraduate degree, she plans to pursue a PhD with a focus on neuropharmacology. Her work is motivated by a desire to understand and improve the human condition through the study of how biologically active substances, including added sugars, interact with neural and physiological systems.

Abstract

When one becomes conscious of hunger and begins to seek food, there are two pathways underlying the feeding mechanism: homeostatic and hedonic. Homeostatic feeding is consumption of food to maintain metabolic function while hedonic feeding is the consumption of palatable foods. Palatable foods often involve the presence of refined sugar, a substance that has been processed and removed of micronutrients to be in its most concentrated form. This includes sucrose and high fructose corn syrup.

Both pathways involve the use of the mesolimbic dopamine pathway which is responsible for motivation. Consumption of refined sugar particularly triggers this pathway due to the release of dopamine from the activation of the Nucleus Accumbens. This is perceived as enjoyable sensations associated with the taste of sugar. Additionally, during homeostatic feeding, consumption stops when satiety is reached, but when refined sugar is consumed, it often overrides satiety leading to overeating and excess sugar consumption.

This overconsumption in turn can activate the addiction pathway due to the positive reinforcement of consuming sugar and the goal to avoid adverse effects as supported by the Opponent Process Theory of Addiction. This theory states the initial use of a drug is for a positive emotional response while continued use is to reduce aversive symptoms of withdrawal and tolerance. The addiction cycle includes binging, withdrawal, and anticipation. Because of the overlap between activation of the mesolimbic dopamine pathway by refined sugar consumption and its role in addiction mechanisms, habitual consumption of added sugar demonstrates addiction-like behaviors. The long-term consumption of this sugar leads to various health conditions including Type II diabetes, cardiovascular disease, and obesity.

Significance of Topic & Broader Impacts

As of 2022, 1 in 8 people were living with obesity, 43% of adults aged 18 years and over were overweight, 16% were living with obesity. Obesity has been tied to various health consequences such as cardiovascular diseases, diabetes, cancers, neurological disorders, chronic respiratory diseases, and digestive disorders. Being overweight in adolescence has also been tied to an earlier onset of these diseases and disorders as well as secondary social consequences such as lower performance in school and bullying.

High body weights and obesity are a result of imbalance between energy intake and energy expenditure. While obesity is a complex problem and multifactorial disease, the inclusion of refined sugar in various forms of food has been tied to obesity rates globally. Sugar is a calorically dense food meaning small amounts can compose the entirety of the daily recommended value of no more than 10% of calories for a 2,000-calorie diet. Due to the wide availability of palatable foods that are composed of dense amounts of refined sugar in the western diet, it can be easy to exceed the recommended amount of daily added sugar consumption. Because of its increased consumption, research in t he habitual consumption of refined sugar is extremely important. While glucose is needed to maintain metabolic functions, artificial forms such as sucrose may take advantage of the neurological and biological processes that ensure the maintenance of homeostasis. Additionally, consumption of palatable foods triggers similar mechanisms in the mesolimbic dopamine pathway as habitual substance use that leads to addiction. Due to this, it is important to better understand how repeated consumption of refined sugar in excess amounts can change neurological systems.

Introduction

In western societies, diets are often criticized due to high amounts of added sugar. In the United States alone, two-thirds of the population are either overweight or obese. Hedonic feeding often manifests in the form of consuming palatable foods that are high in fats and sugars. Simultaneously, when looking at the average ingredient label, it can often be shocking to notice the many forms refined sugar may appear, including dextrose, fructose, barley malt syrup, and more. The two most common forms of refined sugar are sucrose (table sugar) and high fructose corn syrup. In the body, disaccharides and monosaccharides are digested at different rates as the body needs to break disaccharides down into monosaccharides.

When examining obesity levels, there is a positive correlation between obesity and sugar consumption due to sucrose rich foods often being calorically dense. Biological and neurological mechanisms within the body point to the cycle of addiction, which includes binging, withdrawal and anticipation forming around habitual sugar consumption. The mechanisms of this cycle involve the mesolimbic dopamine pathway, which is responsible for motivation and reinforcement of habitual activity. Additionally, the Nucleus Accumbens within this pathway experiences increased activity during the binging stage of the cycle. The relationship between hunger and its regulation by glucose seeking moreover involves the mesolimbic dopamine pathway and its neurotransmitters of dopamine, opioids, and glutamate.

This review serves to develop an understanding of the relationship between habitual sugar consumption and addiction. This will occur through forming a definition of what refined sugar is, how it is digested in the body, and how changes in the nervous system caused by habitual consumption of these sugars mirrors the changes caused by habitual consumption of addictive drugs. Finally, the implications including health issues that arise from excessive sugar consumption such as cardiovascular disease, diabetes, and obesity will be explored.

Defining Refined Sugar

Refined sugar can often be found in various kinds of foods; this includes foods notorious for their sweet flavor profile as well as foods where inclusion may seem less apparent, such as certain condiments and pasta sauce. Refined sugar differs from natural sugar due to its processing, which strips away micronutrients; and isolates it into its purest forms. Refined sugar includes glucose and artificially added fructose such as high fructose corn syrup (HFCS). This form of sugar differs in structure when compared to natural sugar including fructose, glucose, and lactose. This is due to the composition of bonds that determine whether the form of each sugar is a disaccharide or monosaccharide. Glucose appears in the body as the body's preferred carb-based energy source and additionally serves as the body's building block of carbohydrates.

Appetite Regulation

Feeding behaviors can be classified into homeostatic feeding to maintain metabolic function and hedonic feeding when consuming palatable foods. When the sensation of hunger is perceived, there is a disturbance to glucose homeostasis. This is the traditional mechanism that fuels homeostatic feeding. From this, there are several mechanisms triggered within the brain and the gastrointestinal (GI) track. First, orexin levels rise in the brain, promoting meal initiation. This activates first order neurons in the arcuate nucleus of the hypothalamus, including the Aguti-related protein (AgRP) neurons, which co-express neuropeptide Y (NPY) that project to second order neurons in the Lateral Hypothalamus (LHA) and Peri-fornical area (PFA). Stimulation of AgRP and NPY cause release of hormones released in the GI tract. Ghrelin is the only hormone found here that promotes food intake. The trigger of the release of ghrelin is an empty stomach. Satisfaction between and during meals is also signaled by the GI tract. Within meal satiety is caused by cholecystokinin (CCK) from the small intestine and amylin from pancreas, while between-meal satiety is signaled by Peptide YY (PYY) and Glucagon-like Peptide 1 (GLP-1) from the large intestine. Additionally, longer term energy balance is regulated by insulin. Insulin levels in the blood rise when blood glucose levels are decreased. Without insulin, glucose cannot enter the cell. A lack of leptin may also increase appetite by regulating the activation of first-order neurons in the hypothalamus. When leptin is present, it removes inhibition from appetite-stimulating NPY and AgRP neurons in the arcuate nucleus, tipping hypothalamic signaling toward hunger.

Simultaneously, the mechanisms behind the motivation and enjoyment of eating involve the mesolimbic pathway. This motivation is influenced by several environmental cues that utilize sensory systems such as olfactory and visual systems. When consuming a palatable food, there is often a sensation of satisfaction, liking, and enjoyment while foods that are aversive such as bitter food may trigger unenjoyable sensations such as discomfort often indicating consumption of a harmful substance. Activation within the mesolimbic pathway occurs during hedonic feeding. It has been found in rats that during hedonic feeding, when sucrose is consumed, there is an increase in dopamine activation in the Nucleus Accumbens thus demonstrating the link between palatable foods and an enjoyable sensation. In this study, microdialysis and sham feeding were used to monitor the concentration of dopamine in the Nucleus Accumbens.

Moreover, during homeostatic feeding, rat models will continue to feed until satiety is reached. This is triggered once food reaches the small intestine by the release of CCK. In rat models where a chronic gastric fistula was created within the rumen of the stomach, this satiety was never reached. However, during hedonic feeding, rats will continue to feed despite reaching satiety. This is demonstrated through the hedonic feeding of cue-conditioned rats, who consume larger amounts of food when both sated and hungry compared to control rats. These studies demonstrate that a learned cue can contribute to the increased consumption of palatable foods. However, this researcher may not match normal rat satiety behaviors due to the artificial opening caused by the gastric fistula which alters normal physiology of feeding behaviors.

Refined Sugar Digestion and Appetite Influence

Refined sugar often takes longer to digest compared to natural sugar. This is because in natural sugar, there is a presence of micronutrients that often serve as buffers, leading to slower digestion. Because of this, when refined sugars are consumed, there is a spike in blood sugar. Repeated spikes in blood sugar can lead to insulin resistance, which impairs the ability of cells to take in glucose and use it for energy. For example, Wistar rats, who consumed 40% sucrose solutions for a span of 180 days, showed early signs of glucose intolerance and insulin resistance. Additionally, in healthy men between the ages of 19 to 25, who consumed beverages enriched with various amounts and forms of refined sugar, there was an increase in insulin resistance. Insulin plays an important role in hunger regulation when it allows glucose to enter the cell. This occurrence of glucose entering the cell often indicates satiety. A loss of insulin signaling in NPY neurons, which may occur during an increase in insulin resistance, led to increased food intake in mice.

Leptin is a hormone released by adipose tissue. It helps maintain weight by playing a critical role in satiety in the body. When leptin is released, it binds to the LepRb receptor in the arcuate nucleus of the hypothalamus.

When rat models were given a habitual fructose diet, there was a decrease in leptin levels and expression of leptin receptors. In a separate study, rats were given fructose-rich (~60% of calories from fructose) drinking water for several weeks and demonstrated leptin resistance. In these studies, there is however, a lack of diversification in the kind of refined sugars present. While both studies did come to a similar conclusion, they both utilized fructose and no other forms such as sucrose.

Addiction and Craving Mechanisms

Addiction is a state of psychological and/or physical dependence on the use of drugs or other substances, such as alcohol, or on activities or behaviors Conceptual interpretation of drug addiction often involves elements of compulsivity and impulsivity therefore, the addiction cycle is broken into three stages. These stages include binging, withdrawal, and anticipation. In the binging state, individuals often experience a surge in dopamine leading to positive and pleasurable sensations such as euphoria. This is due to stimulation of the nucleus accumbens. Brain imaging studies in humans show activation of dopamine and opioid neurotransmitters during consumption of alcohol and other substances. In the withdrawal state, the brain has adapted to repeated substance use, leading to a struggle to maintain normal dopamine levels. This causes withdrawal symptoms such as cravings, anxiety and irritability. In the anticipation stage, there are compulsive thoughts caused by altered activity in the prefrontal cortex.

This cycle was demonstrated in a study using rats as an animal model. In this study, rats were deprived of food for 12 hours and then provided chow and a sugar solution for 12 hours. After repeated exposure, the rats drank the sugar solution aggressively, resulting in a "bingeing behavior." They also demonstrated withdrawal-like symptoms during 12-hour spans without the sugar solution as well as behavioral changes associated with the anticipation stage.

As an individual navigates between the stages in the cycle, there is a shift from impulsivity to compulsivity driven by a shift from positive reinforcement to negative reinforcement. This is supported by the Opponent Process Theory of addiction, which states that addiction is driven by two sides of motivation. The first consumption of a substance is motivated by the positive hedonic response, however, after repeated consumption of a substance for an extended period, there is a shift to consuming the substance to avoid the adverse affective state. To further explore this, rats were placed on a runway with the ability to self-administer cocaine. Initially rats would run quickly down the pathway, however, over repeated trials they began to slow down and show signs of hesitation. This drug became aversive after its rewarding effects faded. However, the methodology constrains the generalizability because the use of the animal model of rats. This is due to the different neurocircuitry and behavioral responses when compared to humans. Additionally, there is a limited behavioral realism due to the lack of human-like exposure of cocaine in the methodology used.

Crossover Between Appetite Regulations & Addiction Pathways

The neural mechanism behind addiction is rooted in the mesolimbic dopaminergic pathway. However, this system is flawed in the sense that it is exploited by artificial stimuli, including refined sugar. In one an experiment monkeys were given unexpected awards, predicted awards, or reward omissions. Electrophysiological recordings were used to record neuron activity. The source concluded by explaining that when pleasurable behavior is engaged in, such as glucose consumption to satisfy disturbances in glucose homeostasis, dopamine is released. Thus the high portion of glucose found in disaccharides often leads to excess dopamine release, leading to overactivation and desensitization. This in turn leads to a reduction in natural dopamine production and sensitivity of receptors. This is the central cause of addiction as decreased sensitivity leads to an increased craving for refined sugar.

When investigating the effects of consuming palatable foods during hedonic feeding, there is an increase in dopamine in the Nucleus Accumbens as well as increased activity in the mesolimbic dopaminergic pathway. Similarly, well known addictive drugs are known to have effects that release dopamine and opioids. This also occurs through the ventral tegmental area (VTA) and the Nucleus Accumbens within the mesolimbic dopamine pathway.

Implications

Consuming naturally occurring sugar, such as mono-

saccharides in the form of glucose and disaccharides in the form of fructose, often does not lead to health issues due to the pairing of this sugar with minerals, antioxidants, and fibers. However, high refined sugar consumption has been tied to various health conditions. This is because this sugar is often not paired with the various components that interact with naturally occurring sugar. Therefore, there is a relationship between increased added sugar intake and cardiovascular diseases. In a 15-year longitudinal study, it was found that those who had a diet consisting of higher amounts of added sugar (17% to 21% compared to 8%) had a 38% higher risk of developing and dying from cardiovascular disease. This has been linked to the consumption of refined sugar raising blood pressure, thus causing chronic inflammation.

Additionally, evidence also points to high consumption of added sugar causing obesity. This is due to several reasons, including the fact that foods high in added sugar tend to be calorically dense leading to higher amounts of calories consumed. This can also be caused by a disturbance in leptin. High consumption of added sugar can lead to a dysregulation of leptin which regulates signals for hunger. This may also lead to insulin resistance, as the body responds to increased amounts of sugar by releasing insulin.

Increased refined sugar intake can also lead to an increased risk of type II diabetes. The insulin resistance defined above can escalate into this form of diabetes. Type II diabetes is characterized by the body being unable to use insulin effectively or not being able to produce enough insulin to control blood sugar levels.

Conclusions

Across behavioral, neurochemical, and neuroimaging studies, a consistent pattern emerges: habitual sugar consumption can produce addiction-like processes through its effects on dopamine, insulin, leptin, and other appetite-regulating signals. This pattern holds across species and methodological approaches, suggesting that repeated exposure to refined sugars alters metabolic hormones and engages mesolimbic reward circuits in ways that promote continued consumption even after physiological satiety is reached. These findings collectively highlight a close relationship between hedonic feeding for pleasure and the recruitment of the mesolimbic dopamine pathway, forming a cyclical pattern of craving, withdrawal-like responses, and reinforcement that resembles traditional models of addiction.

Nevertheless, the existing literature has important limitations. Many studies rely on self-administration tasks conducted in highly controlled environments that lack competing stimuli, potentially inflating the apparent motivational value of sugar. Cross-study comparisons are further complicated by genetic variability among subjects and by the difficulty of translating behavioral responses from animal models to the complexity of human food-seeking behavior. Additionally, several studies struggle to disentangle hedonic sugar consumption from the body's genuine homeostatic need for glucose, obscuring the boundary between adaptive feeding and addiction-like processes.

Given the widespread availability and diversity of refined sugar, including substances such as maple syrup, molasses, and agave nectar, as well as the increasing use of artificial sweeteners that also stimulate dopamine release in the nucleus accumbens, future research should work to identify which specific sugars or sweeteners most strongly elicit addictive or addiction-like responses. Further studies are also needed to clarify whether artificial sweeteners and traditional refined sugars differ in their capacity to activate reward pathways, alter metabolic hormones, or produce withdrawal-like behaviors. Such work would significantly advance our understanding of the mechanisms underlying sugar-related addiction and its implications for public health.

Access references for this paper using the QR code below.


CORKHOWARD
CONSTRUCTION
WE BUILD TRUST
2121 NEW MARKET PARKWAY
SUITE 118, MARIETTA, GA 30067
770.690.0800 | CORKHOWARD.COM


Johnson Controls
Building Technologies & Solutions
Johnson Controls' Building Technologies & Solutions is making the world safer, smarter and more sustainable, one building at a time. Our technology portfolio integrates every aspect of a building – whether security systems, energy management, fire suppression or HVACR – to ensure that we exceed customer expectations at all times. We operate in more than 150 countries throughout our unmatched network of branches and distribution channels, helping building owners, operators, engineers and contractors enhance the full lifecycle of any facility. Our arsenal of brands includes some of the most trusted names in the industry, such as Tyco®, YORK®, Metasys®, Ruskin®, Frick®, PENN®, Sabroe®, Simplex® and Grinnell®.
For more information, visit www.johnsoncontrols.com
or follow @JCI_Buildings on Twitter.
Atlanta Office: (770) 663-0663
Service Dispatch: (866) 635-1395
Johnson Controls


"Do not follow where
the path may lead,
go instead where
there is no path
and leave a trail"
-Ralph Waldo Emerson


JH
EQUIPMENT, INC.
J&H Equipment Inc.
A World Leader in Vibrating Wire Screens
Designing and Manufacturing Equipment
for the Process Industries
140 Sunshine Way, Alpharetta, GA 30005
770-992-1606 (Phone)
770-992-1983 (Fax)
www.jhequipment.com

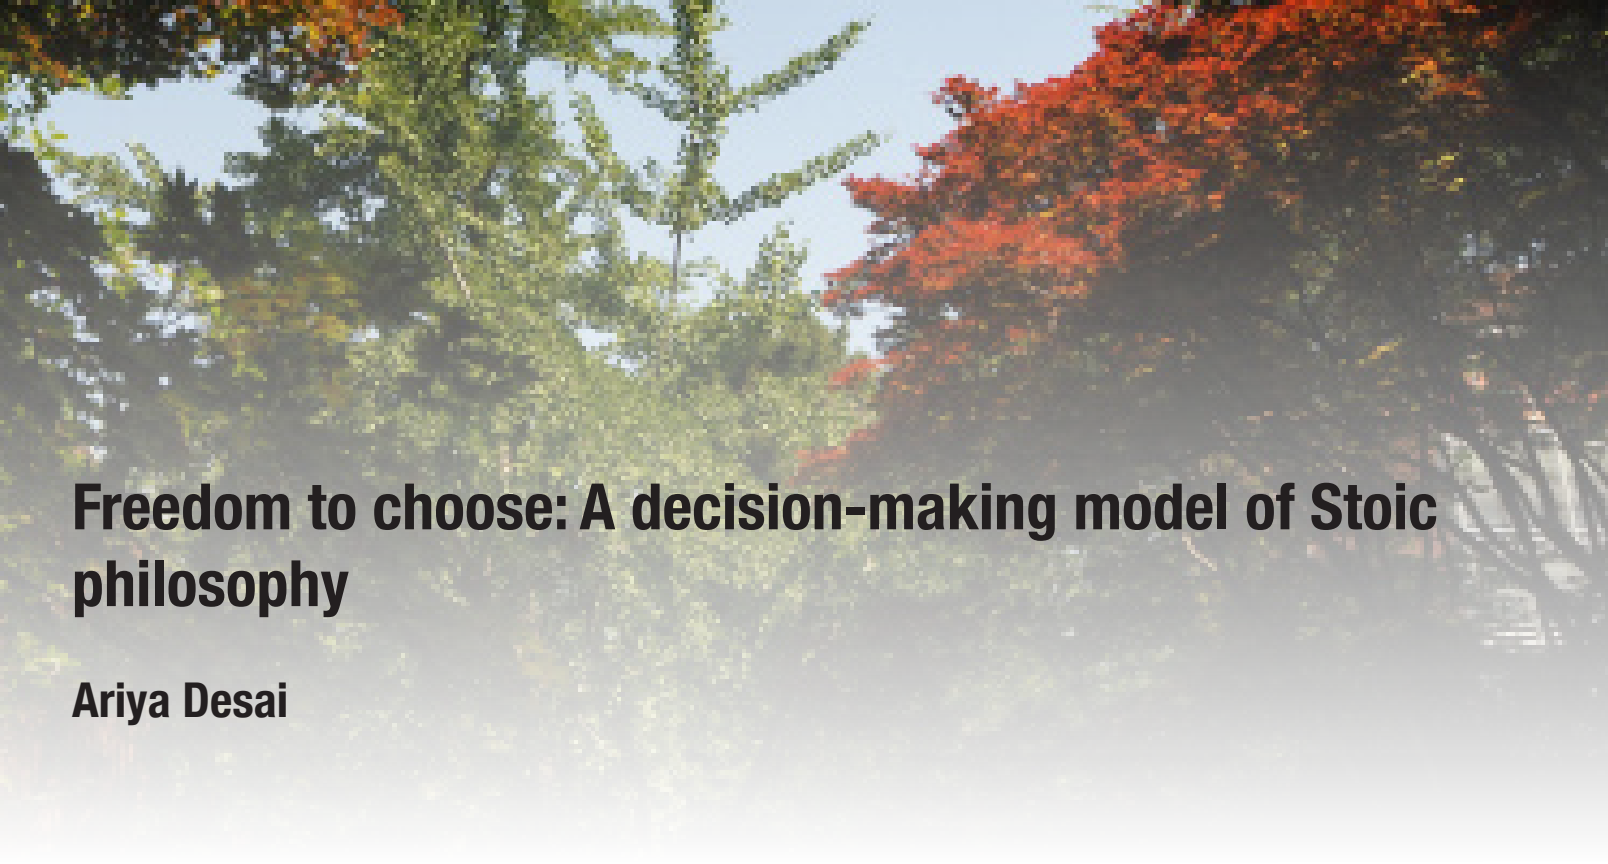

Freedom to choose: A decision-making model of Stoic philosophy

Ariya Desai

Ariya is a senior studying neuroscience at the Georgia Institute of Technology and is working on a minor in Health and Medical Sciences. She is interested in modern applications of Stoic philosophy, particularly in supporting mental health treatments, and explored this topic through her course in the Neuroscience of Behavior. She hopes to use the combination of philosophy and neuroscience in both personal and professional applications in her future in the corporate healthcare field.

Abstract

Stoicism is the ancient moral philosophy focusing on the principles of self-regulation and recognition of what is or is not under an individual's influence. The philosophy is an early antecedent to cognitive behavioral therapy and uses principles of intentional cognition like mindfulness and self-control but has been mischaracterized over time due to alternate translations and interpretations. This paper explores the decision-making basis of Stoic beliefs, illustrates the neural regions and modulators implicated in pursuit of Stoic ideals, and seeks to clarify conflicting viewpoints on the efficacy of therapeutic Stoicism. A Stoic decision-making model relies on three stages of perception, consideration, and self-regulation, which are well-recognized cognitive processes. Overlaying Stoic virtues like courage, wisdom, temperance, and justice upon cognitive behaviors like mindfulness, self-regulation, and perception and awareness show the neuroanatomical underpinnings of the philosophy. Behaviors within these stages originate from the primary regions of the prefrontal cortex (PFC), amygdala, and ventral striatum, and operate through dopaminergic systems to drive decision-making. This review intends to explain the neuroscientific anatomy and cognitive basis of Stoic decision-making and finds that effective understanding and implementation of Stoicism towards educational workplace training, personal study, or therapeutic settings can support feelings of self-efficacy and promote mental health benefits.

Introduction

Philosophy has long posited the question: is human behavior dependent on free will or preconceived fate? Initial arguments seeking to disprove free will have since been overturned by recent studies that indicate free will and volition do impact human cognition.[1] What is popularly conceived of as Stoic philosophy is a weak, extrapolated derivative of its original use; having a 'stiff upper lip' is a poor extension of true Stoic philosophy. Stoicism is defined as the philosophical school of thought emphasizing consideration of personal choices and recognition of internal and external control.[2] Epictetus, Chrysippus, Zeno, and Marcus Aurelius are accredited with works that describe ancient methods of virtuous living - making choices based on reason, examining the origin, context, and meaning of knowledge, and exercising mindfulness. The principles of Stoicism are later adopted and adapted by Locke's empiricism, Descartes' rationalism, and even Diogenes' cynicism, displaying their value across generations of philosophers.

Though our understanding of cognition and higher-level processing has expanded through imaging technology and applied experiments, general zeitgeist suggests more than ever - adults are skeptical of how much power individual choice wields in the face of world events, economic and social instability, and promises of self-determination; only 39% of current Gen Z adults consider themselves as 'thriving'.[3] As the prevalence of mental health disorders like anxiety and depression increase, it is more important than ever to reevaluate the efficacy of common cognitive therapies.[4] Recent scientific explorations ranging from the 2010s have renewed interest in practices like mindfulness and effortful self-regulation.[5] Practices like Cognitive-Behavioral Therapy (CBT), which are based in Stoic practice, becoming prolific in mental health treatment show a demand for the application of Stoic theory in therapeutic cognitive settings.[6] However, there is a great deal of debate on the true efficacy of Stoicism, with misunderstandings or misapplications of the text driving conflicting perspectives. This paper will present both arguments and define how reasonable application of Stoic decision-making provides more benefits than deficits for individuals with mental health conditions that can be aided by cognitive therapy.

This paper seeks to propose that the philosophical position of Stoicism, when placed into cognitive contexts, can provide a tangible, effective model for intentional, mindful perception and interactions with an individual's environment. This paper's goal is to develop a model of Stoic decision-making which can be modified and applied in individuals to promote general wellbeing, and in limited contexts - to act as a behavioral aid for those managing mild mental health conditions. It is not intended to provide a psychological view of Stoic philosophy's validity, but to illustrate the neuroscientific pathway and regions involved and show how these regions modulate the cognitive behavior of decision-making.

The Perception Stage

The primary assumption of Stoic philosophy states that there is a segmentation between an individual's realm of influence and what is outside of that. Cicero's Paradoxa Stoicorum notes that the end goal of self-regulation is the recognition of one's own facilities for reason, hope, expectations, and perception.[7] This initial stage within the model sets up the expectation that free will enables human beings to assert agency and describes how perception impacts decision-making tasks. It also contextualizes flaws in decision-making task designs based on early free-will designs. Free will is typically categorized as: having alternatives to choose from, control over one's own choices, and actions guided or explained by reason as opposed to chance.[8] Without free-will, we cannot explain the cognition behind long-term decision-making. It is first necessary to establish this precedent that free will is subjectively experienced by people to understand how it influences the behavioral outcome of cognitive processing.

Libet Studies and Free Will

Stoicism's core tenets rely on the central ability of people to make decisions which reflect the quality of their character and manage the resulting consequences. Epictetus' discourses feature heavy conversation on the necessity of man to conduct their choices based on reason[9]; Marcus Aurelius' journal Meditations included reminders that a person's life and the quality of it is determined by the choices they make.[10] Thus when examining the neuroscientific basis of enacting Stoicism - we start with the question of free choice and free will. A decision-making model rests on the foundation that people hold the abilities of volition and self-determination. This premise dictates the

effectiveness of the following cognitive processes: analysis of a situation, recall of previous experiences, inventory of current environment and stimuli, and a final decision.

Libet's initial studies in 1983 were foundational efforts in the field and seemed to answer the question of free will with scientific results: the presence of brain activity in the form of electroencephalogram (EEG) readings before a tapping task was performed indicated that actions depended on prior prescience .[11] The latency timing of these EEG recordings and the readiness potential were measured and analyzed which suggests that a meaningful decision was conducted before motor output was performed.[11] For many years, this study's findings dictated the neuroscience of free will - namely that it could not exist.[12] The Libet format for free-will has long been contended, but neuroscientific evidence suggests that free-will does exist and is illustrated by electrophysiology latency studies. Scientists deconstructed the study's design and pointed out that the participants had been asked to conduct decision-making tapping tasks with spontaneous decision-making. These scientists highlighted that the readiness potential observed indicated the decision-making process, not the conclusion.[1,13] These original Libet studies have been questioned on account of the fact that spontaneous reactions to stimuli cannot account for long-term reward evaluation and memory recall which influences long-term, practical decisions in day to day life.[12] Qualitative analyses of the Libet paradigm indicated the variation from study to study and participant to participant might allow conclusions about temporal judgement, but a lack of conclusive evidence on the existence of free will.[14]

It is necessary to establish free will as a core tenet of Stoicism; free will determines a person's belief in their ability for perception of choices and environment to matter.[12] The capacity for free will is what enables sensory perception of external stimuli to occur and also what enables this perceptual information to be considered or discarded in favor of internal mentation and mental causation.[15] Two potential explanations of computational neural modeling can be used to explain free will: the presence of random stochastic noise or the integration of different sources of contextual evidence.[16,17] Schurger et al. examined if the decision-making process in the Libet task was the culmination of neuronal noise rising to cross the threshold; this sort of integration-to-bound model was shown to explain the timing of the readiness potential as a drift-diffusion model showed good fit when applied to the evidence accumulation process.[16] This experiment and model would explain that rather than a true decision occurring milliseconds before a report of conscious decision making, the decision occurs very close to the moment when awareness of the decision is vocalized. Further analyzing alternative computational models to explain free will, a particular subset of neurons in the posterior parietal cortex (PPC) were shown to process and drive history-dependent choices.[17] By using a visually instructed action selection task for mice and conducting two-photon calcium imagery of the PPC neurons, the researchers established that there was significant choice involved in the behavioral trial outside of stochastic noise, and that the PPC neurons play a necessary role in planning and executing decisions, as inactivating these neurons affected history-dependency of choices.[17] These experiments add significant evidence that the subjective experience of free will can be corroborated by effective mathematical modeling and the isolation of neural substrates which account for the kind of planning and choice that is observed during Libet-type tasks. By further experimenting on Libel's constructs and the typical Libet tasks, we can observe that the original experiment's evidence on the early timing of the readiness potential does not disprove free will.

Perception and Decision-making

Having presented evidence supporting that free will exists and humans are agents of their own volition, we circle back to the importance of perception, which is the groundwork for the Consideration and Decision stages. This would pertain to how we examine external stimuli and environments before a decision is made. Per Epictetus' Enchiridion, the first step towards developing self-regulation is for an individual to examine their circumstances and be able to distinguish what is or is not up to us - only then can individuals apply the process of decision-making in an attempt to change the circumstances they do have control over.[9] Human perception of a decision's outcomes relies on the input variables which will impact a decision - task parameters, the physical environment, our personal body's condition and health, and internal cues such as hunger, thirst, or instinctual behaviors. Perceptual

information related to decision-making is encoded across a network of areas which combine the salience network and dorsal attention network - the frontal eye field, insula, intraparietal sulcus, posterior cingulate cortex, and superior frontal gyrus – which control the overt direction and transition between objects of attention.[18] The integration of these regions in decision-making informs future decision outcomes by helping a person to manage goal-directed information and is integral to ensure the downstream assessment of risk and reward.

Primary literature on decision-making concentrates on evaluating performance on immediate tasks or on the computational model of decision-making processes.[20] This format fails to consider the long-term variables which impact decision-making tasks. The congruence effect was found active in a perceptual detection task when a high probability cue induced higher confidence for a target-is-present response, showing that prior knowledge of probability, sensory info, and response outcomes were all amalgamated to provide a confidence judgement.[21] As such, current decision-making tasks are insufficient in capturing the involvement of variables such as increased decision options, prior experiences, delayed feedback, social interaction and influence of feedback, interdependencies in choices, and potentially biased vicarious experiences.[22] Modification of tasks is necessary to include the various feedback mechanisms which cause individuals to reframe, alter, or hesitate in the process of decision-making. The cognitive behavioral field's understanding and application of Stoic decision-making can be improved by deploying delayed feedback, social interactions, and complex choices in decision-making tasks.[22] Most critical decision-making tasks in studies tend to focus on spontaneous or evolutionary reward-driven behaviors, which can be insightful in substance-related or sensation-seeking scenarios but have limited application beyond experimental settings. By developing decision-making tasks that factor in immediate and long-term variables - like sensory stimuli and previous experiences - we can effectively integrate Stoic perceptual theory into cognitive decision-making.

Stoic reasoning advises that human perception of the world should be unbiased, not allowing personal emotions like anxiety, stress, jealousy, fear, and excessive joy or satisfaction to color our objective understandings.[9] It is important to highlight that there is a discrepancy between ideal Stoic teachings and the cognitive load of past experiences; historic experiences and judgements are by nature subjective and unique to persons. The process of decision-making factors of lived experiences, emotions, and reward evaluation, and considerations of feasibility may drive decisions in a way which is rational or acceptable to the self, but not to someone else. This is one limitation of Stoic principles which does not impede the model but does act as a ceiling: while we can ground Stoic perception in cognitive decision-making, Stoic 'philosophy' is an idealized and utopian vision of regulatory behavior and perception with the external world, and effective decision-making will always be colored by subjective perspectives. Using a theory like CBT in mental-health treatment can build the skill of reducing the impact of subjectivity on perception of scenarios and can help to modify and self-correct these biased perspectives towards a more objective decision-making framework.[23]

The Consideration Stage

The Consideration Stage is where the process of 'mentation' occurs - the process of deciding, as opposed to the final concrete decision. The prefrontal cortex, amygdala, hippocampus, thalamus, and anterior cingulate cortex have all been implicated in the early stages of decision-making, from value evaluation to intentional determination of actions.[24]

Reward and Value Judgement

Stoic considerations of reward are two-fold: firstly, that to consider that one's own choices and perspective can determine the level of fulfillment one feels, and second, that the ability to construct our own meaning and reward from challenges is the unique benefit of human cognition.[9] Stoic definitions of virtue do not always coincide with what may be summarized as animal or inherent reward: access to excessive resources guarantees animal survival, but Stoic theory advises that immoderate human attachment to material wealth is improper and should be discarded. So how can we connect the idealistic terms of Stoic value judgments with observed experimental reward analysis? In the Consideration Stage, the solution is to emphasize goal-directed behaviors and tasks in the context of higher

cognitive neural pathways as opposed to lower-level instinctual reward-drives. Perceptual decision-making tasks indicate that the left dorsolateral prefrontal cortex (ldlPFC) modulates abstract decision-making and task preparation[25], while complex decision-making tasks recruit areas associated with internal needs (hypothalamus, midbrain), memory (hippocampus), and relay functions (cerebellum) to support reward representation and evaluation.[26] This second category of neural regions differentially activate and communicate in order to first establish sophisticated reward values and then judge which option is most beneficial.

Beyond judging objective benefit to an individual, what differentiates a 'Stoic' thinker from other schools of philosophy is the focus on tailoring their choices to meet certain ideals for behavior. These ideals are defined as the core virtues of Stoicism and include abstract aspects of morality like courage or wisdom.[9,10] These virtues inhabit subjective definitions from individual to individual. The challenge with Stoic applications in cognitive reward judgement is that probing the circuitry of reward judgements is based in rodent-model studies involving time-sensitive reward-punishment decision tasks. In an example of rodent studies testing risk aversion judgements, sucrose (rewarders) or air-puffs (punishers) are applied to a rat to judge which stimulus will produce increased aversion - whether the rats stayed or moved away (Neville et al., 2020). A computational analysis of results found that the rats showed higher weighting of punishment stimuli versus reward stimuli, and that rats who were rewarded with sucrose before the test were more averse.[27] This experiment highlights how researchers can study decision-making in mammals based on their expectations of a decision's outcome. Furthermore, organisms are shown to typically choose distant rewards only if it is a shorter distance of time in the future and if the value derived is large relative to a closer reward.[28]

While these experimental set-ups are useful to measure short-term judgement in mammalian models, human stressors tend to inhabit a different niche and are often long-term - choices around lifestyle, career, family, and futures tend to be more complex than the potential for disturbance via puffs of air or the potential of electric shock. To further elucidate the question of long-term reward analysis and provide future directions, it is worth looking into the putamen neurons, which were shown to change shape from clusters to rectangular geometry as learned adaptive value guided behavior occurs, showing that these neurons are associated with long term decision making and the weighing of values and context over time.[29] Additionally, dopamine is the key neurotransmitter involved in the mesolimbic reward pathway, and in therapeutic settings, modulating dopamine activity may be key in establishing decision thresholds and promoting long-term value judgements.[30]

Contention on Loci of Control

The locus of control theory outlined in popular psychology delineates the difference between individuals who think factors in their life are within or outside of their control: internal or external loci of control, respectively.[31] This social learning theory-based concept has been studied as a factor in core self-evaluation, and how it correlates with an individual's self-efficacy, self-esteem, and factors of wellness like motivation, work performance, and satisfaction.[31] Having an internal locus of control forms a central pillar of Stoic logic - the understanding that a person themself is the only determinant of their mental condition, despite external conditions or hardships. The neurological basis of the loci of control theory is widely contested as it is difficult to localize whole-brain functions like self-perception and abstract conceptualization; instead, it is more helpful to view the application of loci of control onto the cognitive process which is the result of having either locus of control: decision-making.

Typical studies on loci of control focus on survey-based correlation studies without distinguishing how and why an internal locus of control can improve quality of life, wellness and personal care, or feelings of motivation and satisfaction. While less definitive evidence in comparison to experimental studies, correlational studies can initiate the process of narrowing down circuitry related to the locus of control theory. In an aging population, it was found that believing in an internal capacity for growth or change can promote retention of cognitive performance.[32] They found that the neural regions tied to emotional significance (amygdala), cognitive evaluation of emotions versus goals (anterior cingulate

cortex), and top-down control (right middle front gyrus) were implicated in having an intrinsic locus of control by acquiring fMRI imaging during daily tasks like managing finances or doing puzzles.[32] The concept of being able to direct will and attention with intentionality is linked to the third wave of cognitive behavioral therapy and may be valuable in building mindfulness.[33] The third wave of CBT focuses on a shift from changing the feeling or frequency of 'negative' emotions to exploring the way a client or patient feels or relates to that experience.[34] When applied to Stoicism, this method would involve integrating Stoic principles like self-regulation or objective rationalism into how someone processes their emotions or experience mental health condition - patients would need to consider the way that they feel about Stoicism and how applying it may change their experience.

While having an internal locus of control has been linked in previous literature to measurable positive outcomes, there is value in studies measuring how individuals turn this mindset into action in an experimental setting. By measuring dependent variables like motivation and satisfaction, the correlation between having an intrinsic locus of control and applying it to decisions can be statistically validated. When subjects in an adult study on intrinsic motivation were asked to recall instances of self-driven learning, growth, or emotional regulation - the ventromedial prefrontal cortex, which drives reward processing and storage, and the anterior cingulate cortex, which evaluates personal importance and regulates emotional processes in decision-making, showed significant activation.[35] Held together, the activation of these regions represent the concepts of hedonia, or positive emotions, and eudaimonia, or meaning and availability of opportunities.[35] Not only does internal motivation in long-term memory impact future decision-making, but it also holds value by influencing individuals to be able to reflect on the meaning of self-regulation for improved mental health. This upholds the Stoic suggestions by Epictetus and remnants of Zeno of Citium's papers from Diogenes Laertius' translation which advise that events cannot perturb a person if their faculty of perception and reasoned judgement is sound.[2] Following these pieces of evidence, the Consideration Stage within the decision-making model illustrates that we can get closer to the challenging standard the Stoics set by weighing potential rewards in a way that activates the higher cognition pathway mediated by the PFC and hippocampus and by exercising an internal locus of control.

The Decision Stage

The final stage in the model proposes that there is a metacognitive point in the decision-making process where a conclusive decision is made and converted into intention of execution. This concerns the Stoic concept that virtuous or good contemplation needs to be reflected within and proven through virtuous actions; Marcus Aurelius reminds himself in Meditations that it is not enough to think of oneself superiorly as a student of philosophy, but that it must be executed in his duties as an emperor, student, family member, politician, and public servant.[10] In the way that potentiation enables synaptic strengthening at the molecular level of learning - repeated virtuous actions strengthen confidence in decision-making and establish patterns for future judgement calls to reference.

Execution of Free Will

The final critical juncture in the decision-making process occurs at the apex of cognition - the point at which planning, experience, and metacognition are converted into tangible action. Libet's studies saw this action as being the spontaneous culmination of immediate stimuli, but experimentally, this can be defined as the point at which individuals come to a decision and are conscious of what this decision is. Typically, Stoic thinking is often mistaken for excessive self-reliance and strict self-discipline, but effective programming of Stoic principles has been applied in some high-stress industries like military and medical organizations.[36,37] Identification of the pathway of Stoic decision-making execution yields information on which neural regions are involved in decision-making. fMRI studies on intrinsic, self-motivated behaviors as opposed to instinctual actions implicated a large, diffuse network of regions which support goal-directed behavior and executive control, including the middle cingulum, the middle and inferior frontal gyrus, the supplementary motor area, the anterior insula, superior and inferior parietal lobules, and the cerebellum.[38] The authors found through PET/fMRI studies as well as later seed-based fMRI functional connectivity analysis that these locations in the brain form a gradient, where the anterior regions are responsible for intentionality

and posterior regions are tied to the process of monitoring actions and their outcomes.[38] Mapping of decision making cognitive processes yielded that the lateral PFC, amygdala, and orbitofrontal cortex – which conduct the different levels of processing which occurs during decision-making - are involved in the holistic process of decision-making as well as the category of neural cells at each level of organization.[39] As such, these are valuable regions to highlight for future study of Stoic thinking. While we have insight into these regions and their discrete role in decision-making, the field's current understanding of Stoic cognition is highly psychological, and our understanding can be improved by experimentally imaging functional activity of these regions while participants actively apply Stoic thinking. Introducing a controlled and measured study of Stoic cognitive decision-making using long-term and short-term task types measuring direct neurological activity as an output can more securely prove the Stoic decision-making models and remains a future direction of study.

Self-regulation in Action

The notable process of self-regulation is the awareness and control of oneself towards a purposeful directive, and involves the reward and emotional pathways, implicating the striatum and amygdala.[40] Self-regulation is regarded as a key cognitive process which matures in proficiency as people cognitively develop and attain practice with discipline and motivated tasks. It can be regarded as the peak of Stoic self-awareness - not only knowing one's environment and personal condition but being able to manipulate it to achieve something you want. An application of hands-on Stoic training practice for medical students reported growth in metacognition, preparedness, and empathic imagination after training concluded.[37] This use of Stoic training, when integrated into a high-pressure field like medical training, indicates that growth in compassion and self-reflection were cultivated by applying a Stoic mindset. Students' mental health concerns tied to peer pressure, toxic positivity, and fear of missing out in life - all negative mental health states - were shown to contribute to Passion - defined as strong emotional conflict, and students practicing Stoicism were shown to reduce scores of Passion.[41] On the other hand, certain Stoic principles may be upheld in a way that improves self-regulation, but with a negative impact on mental health. Pathak et al. share that Stoic ideals typically signal a lack of help-seeking behavior in health beliefs, especially men.[42] Severe interpretations of Stoic philosophy can exacerbate isolation, self-reliance, and austerity - which uphold the principle of self-regulation that the Decision Stage references, but which can harm overall wellness of individuals. This may be attributed to cultural Stoicism, which has emphasized principles like self-control and emotional moderation to a more extreme level, veering on complete suppression and ignorance of one's own feelings.[12] This proves an ineffective model as total repression of emotions or needs is not sustainable; critique of this hardline approach promotes the use of cognitive training, or cognitive redirection and review of one's own feelings.

The idea of eudaimonia, or happiness derived from control over ourselves and our judgements, may also impede progress in mental health treatment as many mental health conditions cannot be treated without interventions, medication, and a combination of CBT.[43] It can be damaging to assert that Stoic decision-making can fully treat and solve mental health issues in the wider population by achieving 'indifference' to mental health concerns.[43] While Stoic philosophy enables a high level of cognitive self-regulation - it can be damaging if the decision-making process is used to engage in harmful, highly regulatory behaviors to the point of self-punishment or harm.

The variables of self-regulation are also tied closely to mindfulness and intentionality. The element of mindfulness is the awareness of one's own feelings and condition and is the most positive application of the Stoic decision-making process. Self-efficacy, humor, and retrieval of positive memories and emotions were shown to enable a feeling of life satisfaction in a survey of Spanish adults.[44] Self-belief and self-efficacy in younger participants of the sample were also associated with using humor and life satisfaction, showing how self-regulation not only can have effects on cognitive performance and welfare but also boost subjective perceived satisfaction and outlook.[44] Mindfulness in students was further shown to improve and alleviate stress and anxiety and promote feelings of mental wellbeing.[45] These examples illustrate that if reasonable self-regulation is applied in situations where individuals have expert guidance, Stoic decision-making has experimental potential to improve general

wellness and promote feelings of mindfulness and self-efficacy.

Conclusion

The three-stage decision-making model of Stoic philosophy proposed here offers an explanation of how goal-directed and value-evaluating decision making can drive improvements in mindfulness, attention, and intentionality. The key stages of Perception, Consideration, and Decision articulate the neurochemical mechanisms behind the process of decision-making and can be refined per modern sensibilities to increase subjective of satisfaction, self-motivation, and free will in practitioners. Several brain regions and pathways, from the salience network to the mesolimbic reward system, are involved in the stages of Stoic decision-making. Examples of Stoic decision-making can be shown through its integration into therapeutic interventions like CBT. The concepts of Stoicism may prove useful when used to corroborate proven clinical methods such as talk therapy and CBT, where it can be used as a supportive tool to improve mental health in individuals, especially in conditions of decreased severity.

Future directions of study in isolating the impact of Stoic decision-making in neural activity may be useful to very clearly connect Stoic scales and measures established by researchers with neural imaging and electrophysiology outputs. It is also necessary to apply Stoic decision-making frameworks in relevant settings: to enable adoption of Stoic decision-making, it is necessary to use the cognitive framework in tandem with required treatment like medications. A key limitation of Stoic thinking is the tendency to assert self-regulation in a way that can detract from seeking help - to be applied to a wider audience, this viewpoint must be amended to explain and support modern mental health conditions which did not functionally or popularly exist when the philosophy was constructed. With clearly identified cognitive groundwork in neuroscientific theory, this paper asserts that ancient Stoic philosophy can be an instrumental guide for intentional decision-making and act as a source of cognitive therapy for mild to moderate mental health conditions. A revival of Stoic practice can reinstate cultural feelings of autonomy and volition, in an era where key health challenges like sleep, diet, exercise, and social connections can be modulated by oneself. In the face of stressors like digital reliance, isolation, and global economic pressures, individuals can feel empowered to improve their quality of life by believing and acting in a way which promotes resilience, community, and simple, meaningful lifestyle changes.[46]

Acknowledgments

I would like to acknowledge Dr. Mary Holder, Dr. Jasmine Hope, and Kali Singh for their guidance and feedback in the process of creating this review paper.

Access references for this paper using the QR code below.

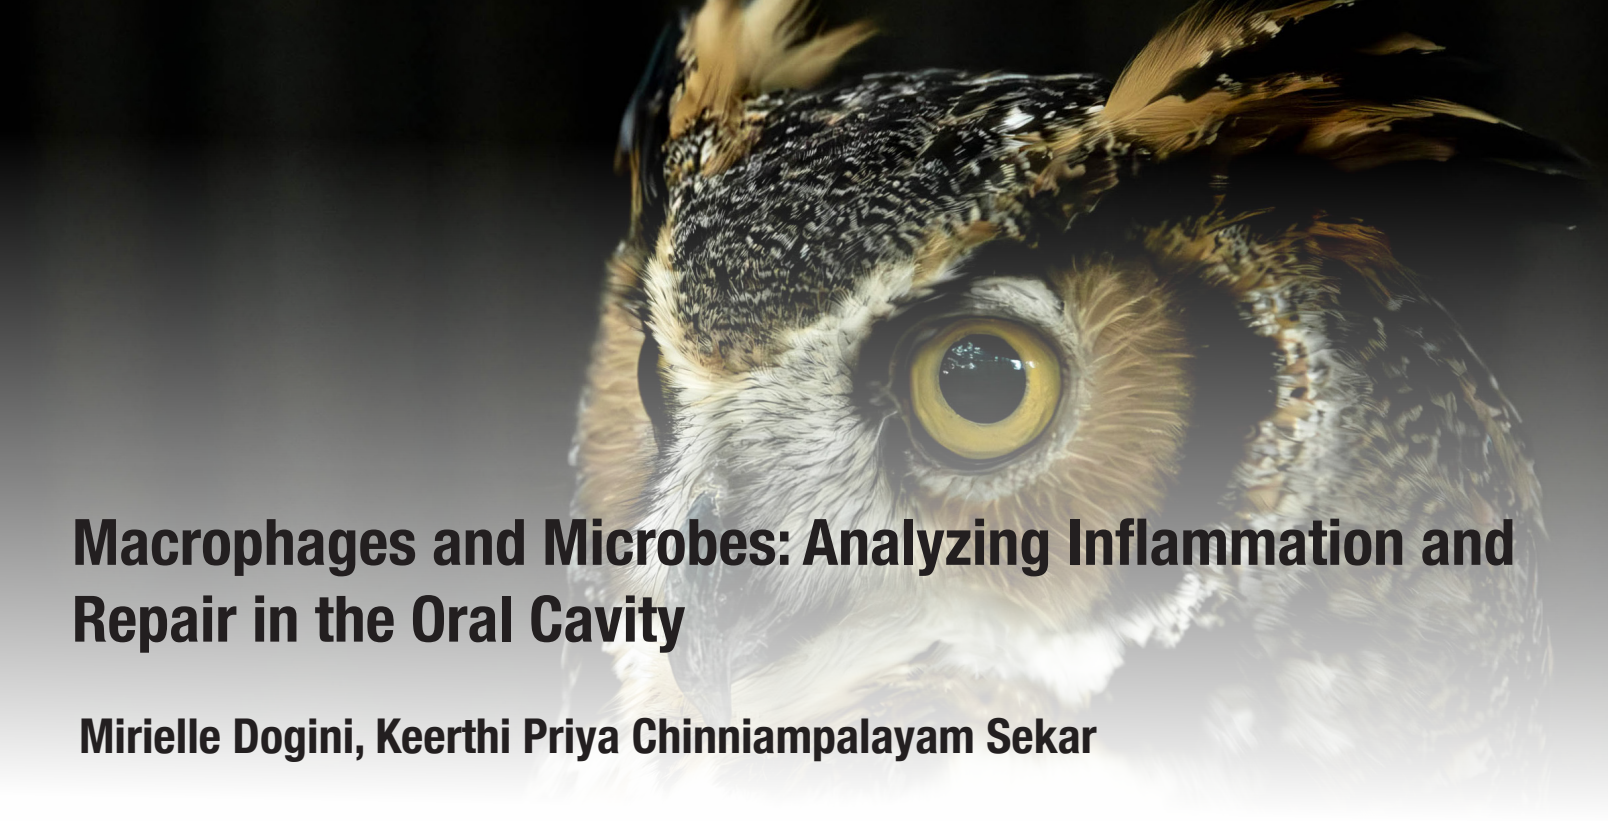

Macrophages and Microbes: Analyzing Inflammation and Repair in the Oral Cavity

Mirielle Dogini, Keerthi Priya Chinniampalayam Sekar

Mirielle was Biology major with a minor in Health and Medical Services. She graduated from Tech in December 2025 and plans to attend medical school. She researched in the Goudy Regenerative Research Lab about cleft palate repair for pediatric patients. She was under the mentorship of Keerthi Priya Chinniampalayam Sekar as Petit Scholar.

Abstract

The oral microbiome exhibits high diversity. It is composed of a community of bacteria, fungi, and viruses that live in the mouth playing a crucial role in oral and systemic health. Considering the oral cavities complexities is essential for understanding recovery of wounds, particularly in clinical contexts such as cleft palate repairs. Cleft palate repair cases can result in impaired healing outcomes such as oronasal fistula. To reduce oronasal fistula formation, strategies such as targeting pro-inflammatory and anti-inflammatory macrophages are being explored to promote immune responses and improve wound healing outcomes. The manuscript explores the impact of added probiotic bacteria and pathogenic bacteria on the use of macrophages to coordinate repair processes through inflammation and tissue regeneration. We hypothesized that probiotic bacteria promote a balanced immune response and accelerated wound healing whereas pathogenic bacteria differ in immune response leading to impaired wound healing. The probiotic strains examined were Lactococcus lactis cremoris (LLC) and Lactobacillus rhamnosus GG (LGG), while the pathogenic species of interest included Porphyromonas gingivalis (PG), Streptococcus mutans (SM), and Enterococcus faecalis (EF). The results indicate that probiotics promote balanced immune activation and improve healing while pathogens produce dysregulated inflammation and impaired healing. Understanding macrophages and oral cavity composition can assist in developing effective treatments across medical disciplines.

Introduction

Oronasal fistula (ONF) is a common postoperative complication associated with cleft palate surgery and is a persistent challenge in craniofacial reconstruction. ONF is described as an abnormal opening in the roof of the mouth between the oral and nasal cavities. Patients often experience air leakage that may cause speech impairments, and lodging of food particles which increases the risk of chronic infection and discomfort [1]. Management of ONF is complex, and closure of the fistula depends on the location, size, tissue availability, prior surgeries, and clinical presentation. Surgical procedures ranging from small local flaps to microvascular tissue transfers have been used for closure of ONF [2]. Despite surgical intervention, recurrence remains a problem with as few as 25% of patients developing recurrent cleft palate fistulas following their initial repair. These complications have a profound impact on the quality of life, contributing to prolonged treatment, speech impairment, and reduced postoperative satisfaction.

New strategies for wound healing are needed to improve ONF development after cleft palate repair. To reduce ONF formation, strategies such as targeted inflammatory and anti-inflammatory macrophages are being explored to promote immune responses and improve wound healing outcomes. Macrophages play a central role in all phases of wound healing due to their ability to promote tissue repair, regulate inflammation and release a variety of inflammatory cytokines, growth factors, and chemokines. Circulating monocytes are initially recruited to the damaged site. During the early stage of inflammation, pro-inflammatory M1 macrophages are involved in phagocytosis and clearance of debris and pathogens [3]. Additionally, these M1macrophages secrete high levels of pro-inflammatory mediators and reactive oxygen species to contain infection [4]. During the later period of inflammation, anti-inflammatory M2 macrophages are differentiated through monocytes or transitions of previous pro-inflammatory macrophages [5]. M2 macrophages support tissue remodeling by promoting angiogenesis, extracellular matrix deposition, and re-epithelization, and support fibroblast activity [6]. Under certain circumstances, macrophages can promote excessive inflammation or fibrosis. Evidence has found that macrophage dysfunction can contribute to the pathogenesis of nonhealing and poor healing outcomes [7]. Thus, a balance between M1 and M2 macrophages is essential. Prolonged M1-driven inflammation can delay healing and impair pathogen clearance and wound stabilization [8], while prolonged M2 macrophage activation can lead to excessive wound healing and fibrosis. Macrophages contribute to wound healing and may be used to identify regenerative properties and improve healing wounds.

Previous research has found that probiotic bacteria may enhance ONF healing by influencing macrophage activity. Scientists found that probiotic strains induce anti-inflammatory cytokine release from macrophages to accelerate wound healing which facilitate resolution of inflammation and promotes tissue regeneration [9]. These immunomodulatory effects can contribute to restoring microbial balance, preventing colonization by pathogenic species, and reinforcing mucosal barrier defenses [10]. Probiotics have shown to have a positive influence on the microbiota and immune system as well as shift macrophage polarization toward the anti-inflammatory M2 phenotype in the gut [11]. This phenotypic shift is essential in transitioning the wound environment from an inflammatory state to one that supports repair. Probiotics can help accelerate wound closure in oral tissues and influence microbial communities' immune responses. Targeting macrophage and microbe interactions using probiotics therapeutics may offer promising applications in cleft palate repair and minimize complications in oronasal fistula. By targeting macrophage microbiome interactions, probiotic-based therapeutics could offer a biological approach to improve healing efficiency and minimize recurrent defects.

Individual variability in wound healing may be influenced by the microbial composition, such as pathogenic bacteria found within the oral cavity. Researchers found that different microbial and pathogenic communities' impact wound healing trajectories [12]. Opportunistic species can interfere with major immune signaling pathways and can contribute to persistent inflammation and delay epithelialization. These species can disrupt host-microbe homeostasis and interfere with key immune signaling pathways essential for tissue repair and leading to infection [13]. Pathogenic species create a dysregulated environment that can further impair fibroblast function, extracellular matrix remodeling, and angiogenesis. Moreover, per-

sistent colonization of pathogenic bacteria can create a chronic inflammatory state where pro-inflammatory M1 profile remains locked rather than transitions to a pro-healing M2 phenotype. This imbalance slows tissue regeneration but also predisposes patients to recurrent fistula formation. Understanding how oral microbial profiles vary between individuals may be important for predicting healing outcomes and the development of target therapeutic strategies that restore microbial balance to support optimal wound repair.

The abundance of pathogenic and probiotic microbial communities in the oral cavity can influence healing trajectories and can explain why some patients experience recurrent ONF while others recover successfully. This can help tailor approaches to therapeutic avenues and determine optimal wound healing based on microbiota composition. It is essential to explore ways to improve wound healing using analysis of cytokine responses and tissue repair. Understanding microbial immune interactions can improve clinical outcomes and inform personalized therapeutic strategies.

We investigated the impact of probiotic and pathogenic bacterial strains on macrophage inflammatory signaling to better understand how microbial composition may influence wound healing trajectories. We focused specially on TNF-α and IL-6 expression and secretion as representative inflammatory mediators involved in early wound responses and immune regulation. By exposing macrophages to cell-free supernatants from probiotic and pathogenic species, we sought to isolate signaling effects and characterize differential cytokine responses. This experimental design allows for comparison of microbial immune responses and provides insights into how oral microbiome composition may shape postoperative healing outcomes.

Methods

Bacterial Inoculation

Probiotic and pathogenic bacterial strains used in this study included *Lactococcus lactis cremoris* (LLC), *Lactobacillus rhamnosus GG* (LGG), *Porphyromonas gingivalis* (PG), *Streptococcus mutans* (SM), and *Enterococcus faecalis* (EF). Each strain was cultured under conditions appropriate for its growth requirements. LLC and LGG were grown in MRS broth at 37°C under aerobic conditions, whereas SM, and EF were cultured in Brain Heart Infusion Broth with EF supplemented with rifampicin; both grown under aerobic conditions. PG was grown in Tryptic Soy Broth and maintained in an anaerobic chamber at 37°C.

Overnight cultures were then standardized for density by measuring optical density at 600 nm. Bacterial suspensions were incubated to reach the mid-logarithmic growth phase. To obtain bacterial supernatant, cultures were centrifuged at 3,000 × g for 10–15 minutes, and the resulting pellet was discarded. The supernatant was carefully transferred to sterile tubes and passed through a filter to remove any remaining bacterial cells. Cell-free bacterial supernatants were aliquoted and stored at −80°C until use. Supernatants were thawed immediately before treating macrophage cultures.

Cell Culture

Murine macrophage cells RAW 264.7 were cultured under sterile conditions in Dulbecco's Modified Eagle Medium supplemented with 10% fetal bovine serum (FBS) and 1% normacin. Cells were maintained at 37°C in a humidified incubator with 5% CO_2 and were passaged at 70–80% confluency using standard aseptic techniques. For experiments, macrophages were seeded in 6-well plates at a density of 3×10^5 cells per well and allowed to adhere overnight. Probiotic and pathogenic bacterial strains were prepared as bacterial supernatants and were added directly to macrophages and incubated for 72 hours. Untreated macrophages were included for baseline comparison.

Quantitative Real-Time PCR

Following bacterial treatment, macrophages were washed twice with sterile phosphate-buffered saline and lysed directly in the culture plate using TRIzol. RNA extraction was performed using chloroform and isopropanol to obtain an RNA pellet. RNA concentration and purity were assessed using a NanoDrop. RNA concentrations were reverse transcribed into cDNA using iTaq Universal One-Step RT-qPCR kit. The resulting cDNA was diluted 1:10 in nuclease-free water.

PCR reactions were prepared using SYBR Green master mix on a Bio-Rad iCycler iQ Real Time PCR Thermal Cycler. Each 20 µL reaction contained master mix, forward and reverse primers, and diluted cDNA. Primer sequences for the target cytokines TNF-α and IL-6, as well as the reference gene, GAPDH. Triplicates were run for each sample. Thermal cycling conditions

were as follows: initial denaturation at 95°C for 2–3 min, followed by 40 cycles of denaturation at 95°C for 10–15 seconds, annealing/extension at 60°C for 30–60 seconds and a final melt curve analysis 65–95°C, 0.5°C increments to verify amplification specificity. Relative gene expression was calculated using the 2^−ΔΔCt method, with the reference gene serving as the internal control [14].

Luminex Cytokine Quantification

Cytokine levels of TNF-α and IL-6 in macrophage culture supernatants were quantified using a MILLIPLEX Panel A magnetic bead–based Luminex assay. Following bacterial treatment, macrophage supernatants samples were centrifuged to remove debris before assay. Samples, recombinant standards, and fluorescently coded magnetic beads were added to 96-well plates and incubated with agitation to allow target cytokines to bind to analyte-specific capture beads. After incubation, plates were washed multiple times to eliminate unbound material and sequentially incubated with detection antibodies and streptavidin-phycoerythrin to generate a quantifiable fluorescent signal. Beads were resuspended in sheath fluid and acquired on a Luminex instrument, with a minimum bead count per analyte collected for each sample. Cytokine concentrations were calculated from standard curves generated using recombinant standards and analyzed with a five-parameter logistic regression model. All samples were run in duplicate, and cytokine concentrations were reported in pg/mL.

Results

To determine how probiotic and pathogenic oral bacteria influence macrophage inflammatory signaling, TNF-α and IL-6 secretion was quantified following ex-

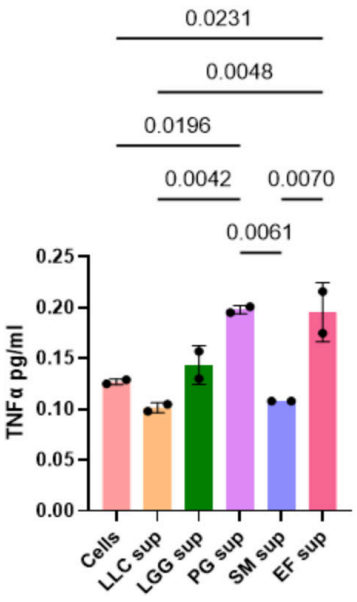

Figure 1. Pathogenic bacterial supernatants increase macrophage TNF-α secretion compared to probiotic strains. RAW 264.7 macrophages were treated with culture supernatants from LLC, LGG, PG, SM, EF. Cytokine concentrations were quantified by Luminex. TNF-α levels remained near baseline in cells treated with LLC, LGG and SM, while PG and EF induced significantly higher TNF-α secretion. Bars represent mean ± SEM. Horizontal lines indicate statistically significant pairwise comparisons ($p < 0.05$). n = 2 biological replicates.

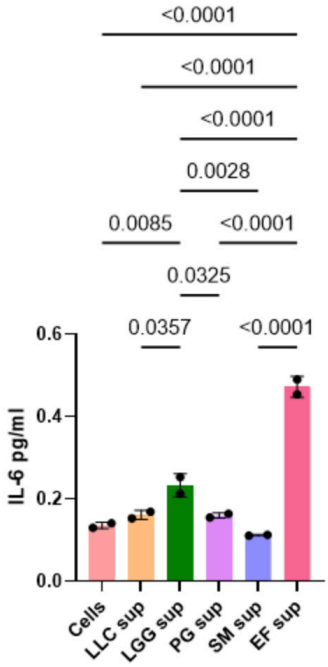

Figure 2. Differential IL-6 secretion in macrophages exposed to probiotic strains and pathogenic bacterial supernatants. RAW 264.7 macrophages were treated with culture supernatants from LLC, LGG, PG, SM, EF. IL-6 responses showed a distinct pattern in which LGG elicited a significant increase compared with untreated cells and exceeded the IL-6 levels produced by PG. EF generated the strongest IL-6 response among all groups. Bars represent mean ± SEM. Horizontal lines indicate statistically significant pairwise comparisons ($p < 0.05$). n = 2 biological replicates.

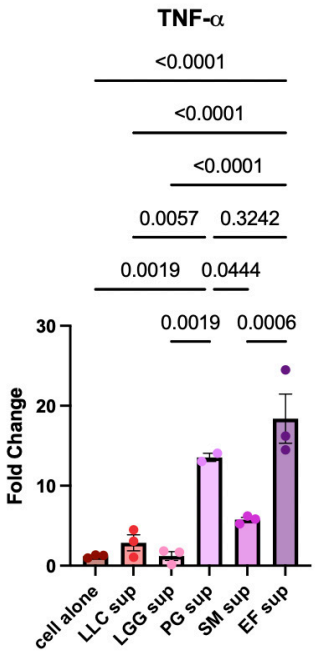

Figure 3. Pathogenic bacterial supernatants markedly upregulate TNF-α gene expression in macrophages. Macrophages were treated with supernatants from the probiotic strains LLC, LGG, PG, SM, and EF. TNF-α transcript levels were quantified by qPCR and are expressed as fold change relative to the reference gene, GAPDH. Probiotic treatments (LLC, LGG) maintained cytokine expression near baseline. Pathogenic bacterial supernatants significantly increased TNF-α with EF producing the highest TNF-α induction. Bars represent mean ± SEM. Horizontal lines denote statistically significant comparisons ($p < 0.05$). n = 3 biological replicates.

posure to bacterial culture supernatants. We discovered distinct cytokine patterns emerged between probiotic and pathogen-treated groups.

Baseline macrophage cultures secreted low levels of TNF-α (~0.12 pg/mL), and treatment with probiotic supernatants from LGG did not significantly alter TNF-α secretion relative to untreated cells. LLC secreted significantly lower levels of TNF-α than the baseline. In contrast, exposure to pathogenic bacterial supernatants resulted in elevated TNF-α production. Supernatants from PG and EF produced significant increases compared with probiotics and untreated cells. Among all treatments, EF and PG elicited the highest TNF-α secretion (~0.20 pg/mL). SM produced lower levels of TNF-α that were not significantly different from the baseline (Fig.1). IL-6 secretion varied across experimental conditions with significance in both probiotics and pathogen. Baseline macrophages and cells treated with LLC and SM had low IL-6 levels and no significant difference. LGG induced higher IL-6 levels than the pathogenic species. The strongest response was observed with EF, which elevated IL-6 secretion to ~0.45-0.50 pg/mL (Fig. 2).

Quantitative PCR analysis revealed distinct transcriptional responses in macrophages exposed to probiotic

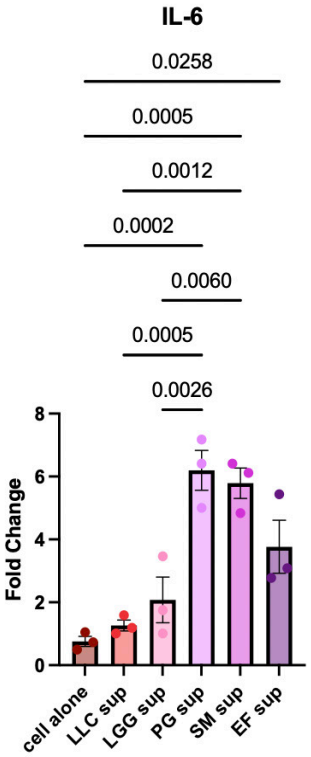

Figure 4. Pathogenic bacterial increase IL-6 transcription relative to probiotic strains. Macrophages were treated with supernatants from the probiotic strains LLC, LGG, PG, SM, and EF. TNF-α transcript levels were quantified by qPCR and are expressed as fold change relative to reference gene, GAPDH. Probiotic treatments (LLC, LGG) maintained cytokine expression near baseline. PG and SM elicited the strongest IL-6 responses, while EF significantly upregulated IL-6 expression. Bars represent mean ± SEM. Horizontal lines indicate statistically significant pairwise comparisons ($p < 0.05$). n = 2 biological replicates.

versus pathogenic bacterial supernatants. TNF-α gene expression remained low in untreated cells and in cells treated with the probiotic supernatants from LLC and LGG, each showing minimal fold-change relative to baseline. PG induced an approximately 10-to-12-fold increase, and SM elicited a comparable but slightly lower response. The greatest induction was observed with EF, which elevated TNF-α expression more than 20-fold relative to untreated macrophages. Multiple comparisons demonstrated significant differences between pathogenic groups and both baseline and probiotic treatments, whereas LLC and LGG did not differ significantly from untreated cells. (Fig. 3). IL-6 expression followed a similar trend. Cells treated with LLC and LGG maintained near-baseline expression. Pathogenic treatments significantly increased IL-6 gene expression. PG produced an increase of approximately 5 to 6-fold, and SM generated a similar magnitude of IL-6 upregulation, with both showing significant differences compared with probiotic-treated and untreated cells. EF produced a smaller IL-6 increase relative to PG and SM but remained significantly elevated over untreated cells (Fig. 4).

Discussion

Wound healing in the oral cavity is defined by microbial exposure, mechanical stress, and rapid tissue turnover which demands a regulated immune response. The study investigated how microbial factors from probiotics and pathogenic species influence macrophage-mediated cytokine signaling. We specifically focused on TNF-α and IL-6 cytokines. Tumor necrosis factor-alpha (TNF-α) is a pro-inflammatory cytokine primarily produced by macrophages during the early phase of immune response, and prolonged expression can impair wound healing by sustaining inflammation and infection [15]. Interleukin-6 (IL-6) is another multifunctional cytokine that can contribute to both pro- and anti-inflammatory responses by helping regulate the transition from inflammation to healing [16]. Balanced TNF-α and IL-6 signaling supports a healthy immune response and tissue repair, while dysregulated and highly elevated levels are linked to impaired healing, fibrosis, and chronic inflammation. Analyzing their secretion and gene expression is critical to understanding differences between probiotic and pathogen-driven immune responses in the oral cavity.
Our results demonstrate that probiotics, LLC and LGG, maintained cytokine levels comparable to untreated cells, while pathogenic bacteria, PG, SM, and EF, induced heightened inflammatory signaling. These findings support the hypothesis that probiotics contribute to immune balance, whereas pathogenic organisms disrupt wound healing through excessive inflammation.

The oral cavity is composed of thousands of microbial species that affect mucosal immunity [17]. When a tissue injury occurs in a cleft palate surgical site, macrophages are recruited to clear debris and modulate downstream healing cascades. In early inflammation, macrophages adopt an activated M1 phenotype to release inflammatory cytokines, phagocytosis of pathogens, and contain infection [18]. Once inflammation is controlled, M2 phenotype promotes angiogenesis, fibroblast activity, extracellular matrix deposition, and tissue remodeling [7]. Excessive inflammation can stall macrophage polarization and impair wound closure.

Our findings align with this immunological framework. Macrophages treated with probiotic bacterial supernatants maintained M1 cytokines at controlled levels which indicates that probiotics do not cause extension of the inflammatory phase. However, pathogenic supernatants elevated TNF-α and IL-6 secretion and gene expression. Particularly, EF triggered the strongest TNF-α responses among the treatment groups. E.faecalis has been found to attenuate proinflammatory cytokine secretions and increase infection and delay healing outcomes [19]. Similarly, P.gingivalis, a key pathogen in periodontitis, is known to induce disease and inflammation as well as prolong cytokine release and skew macrophage polarization [20]. The high inflammatory response we observed in response to PG supernatants reflects these behaviors.

Interestingly, SM showed a more intermediate profile across both cytokines compared to the other pathogenic species. While SM did induce IL-6 gene expression above the baseline, TNF-α secretion and gene expression remained relatively modest. The pathogen may elicit distinct macrophage programs that reflect its unique virulence mechanisms. Additionally, there may be differential macrophage signaling for SM, and this may explain why certain infections by specific pathogens lead to slow-healing fistulas, while others do not.

The probiotic strains demonstrated immunomodulato-

ry effects that are consistent with prior research. LGG slightly increases IL-6 secretion compared to untreated cells but maintains low TNF-α levels. The slight elevation in IL-6 may support wound closure. Since IL-6 has a dual role of initial pro-inflammatory than transition into anti-inflammatory; indicates that LGG may undergo this transition but under a controlled release since TNF-α remained low. Previous research has shown that LGG modulates host immune system and enhance immune defenses against pathogens and reduce inflammation [21]. Meanwhile, LLC appeared to lower TNF-α production below baseline. These responses reflect the finding that probiotics support immune homeostasis and macrophage polarization toward M2 phenotype [22].

Oronasal fistulas remain a clinical challenge, affecting as many as 20-60% of patients depending on anatomical and surgical factors [23]. ONFs can lead to speech dysfunction, and leakage of food and fluids into the nasal cavity. Additionally, it can cause significant emotional burden for young patients and their families as well as reduced satisfaction with surgical outcomes. Given these complications, novel strategies to reduce postoperative wound failure are needed.

This study supports the idea that personalized wound healing strategies could be optimized based on patient's oral microbiome profiles. Individuals with high levels of pathogenic species may enter cleft palate surgery at elevated risk for inflammatory dysregulation. Persistent M1 macrophage activation as shown in the study reflects delayed wound closure and increased fibrosis [18]. However, a microbiome enriched in both commensal and probiotic microbes may support more efficient resolution of inflammation and faster transition to tissue repair.

We demonstrate that probiotics produce substantially lower inflammatory signals, and our findings can highlight potential therapeutic applications. Probiotic supplementation prior to and after cleft surgeries could stabilize wound associated immune response and protect against dysbiosis in the oral cavity. Emerging evidence that probiotics can improve oral wound closure and enhance beneficial immune function as well as promote a balance between oral pathogenic bacteria [24].

The distinction in macrophage responses observed across bacterial groups shows the potential for TNF-α and IL-6 to serve as biomarkers of impaired healing. Elevation in these cytokines may signal a shift toward chronic inflammation that warrants clinical intervention. Therapeutic strategies that limit excessive TNF-α and IL-6 expression, support macrophage transition from M1 to M2, and shift oral microbiome composition toward beneficial taxa could reduce ONF recurrence following repair procedures.

While our experiments focused on inflammatory cytokine signaling, macrophages mediate wound healing through various ways such as recruitment of mesenchymal stem cells, activation of epithelial proliferation, modulation of fibroblast and extracellular matrix interactions, and angiogenesis and hypoxia responses [25]. Understanding the multiple pathways may offer other ways to improve wound healing. Future research should expand profiling to include M2 associated cytokines which would clarify probiotics and pathogen influence on macrophage polarization more precisely. Furthermore, testing both acute and long-term time point would mimic surgical healing progression from days to weeks.

While this experiment used RAW 264.7 mouse macrophages, primary human macrophages may reveal more information on the cytokine profile due to wound healing. Additionally, primary human macrophages may lead to personalized immune signatures important for the clinical context. Incorporating complex 3D wound healing modeling that approximates tissue structure would improve accuracy in predicting in vivo effects.

Despite our findings which support our hypothesis, there are several limitations that are present in the study. We used supernatants rather than live bacteria which limited the full examination of bacterial mechanisms such as adhesion, invasion, or biofilm formation. Additionally, the sample size in the cytokine analysis was small and this can limit statistical power. RAW 264.7 cells have limited complexity and plasticity compared to monocyte-derived macrophages. Another limitation is that cytokine measurements were limited to two inflammatory markers which did not reflect the full cytokine profile of macrophages and reparative functions during wound healing. Lastly, pathogenic responses may differ from real oral environments.

Addressing these factors in future studies will allow stronger conclusions with direct relevance to surgical healing.

Conclusion

This study demonstrates that probiotic and pathogenic oral bacteria exert fundamentally different influences on macrophage-mediated inflammatory responses. Probiotics such as *Lactococcus lactis cremoris* and *Lactobacillus rhamnosus GG* helped maintain cytokine levels near baseline, suggesting that they support a more regulated transition during inflammatory phases of wound healing. In contrast, exposure to supernatants from pathogens such as *Porphyromonas gingivalis, Streptococccus mutans*, and *Enterococcus faecalis*, resulted in dysregulated TNF-α and IL-6 signaling patterns tied to sustained inflammation and delayed tissue repair.

Our findings provide a biological foundation for understanding microbiome therapeutic strategies in cleft palate postoperative care to prevent oronasal fistula development. By investigating the benefits of probiotics and their ability to influence macrophage activity, targeted intervention may improve outcomes for surgical approaches. Ultimately, individualized wound management that considers both oral microbiome and macrophage behaviors has the potential to enhance healing outcomes. Integration of microbiome science with immunology into treatment strategies may reduce recurrent rates, support better recovery, and improve quality of life for patients born with cleft palates.

Acknowledgements

The author would like to acknowledge Keerthi Priya Chinniampalayam Sekar and Steven Goudy for their mentorship and support.

Access references for this paper using the QR code below.

Menstrual Product Sustainability: Are we there yet? A Look into United States and European Union Practices

Marisa Perera[1], Ophelia Winslett[2], Vibha Iyer[3], Kayla Carneal[4], Mallika Digavalli[5]*, Arina Kucinskaja[6], Aline Banboukian[7], PhD., Jasmine Burton[8], MPH, MBA

Marisa Perera is a Biomedical Engineering Master's student at Georgia Tech. She earned her Bachelor's in Biomedical Engineering with a minor in Robotics in 2025 from Georgia Tech. Marisa was the project leader for this paper, and eventual co-director of Wish for Wash, working with the team for four years to improve sustainability in menstrual products and leading the research in plastic usage and alternatives in menstrual products. Outside of Wish for Wash, Marisa researches flexible biosensors with hopes of working in the medical device industry.

Ophelia Winslett is a Mechanical Design Engineer within the Neuroscience division at Medtronic. She earned her Bachelor's in Biomedical Engineering with a minor in Physiology in 2024, followed by her Master's in 2025, both from the Georgia Institute of Technology. A member of Wish for Wash for four years, Ophelia's research involvement focused on evaluating the sustainability and ESG transparency of major U.S. companies within the menstrual hygiene market.

Vibha Iyer is a PhD student in Robotics at Georgia Tech, where she is pursuing her Masters in Mechanical Engineering. She earned her Bachelor's Degree in Biomedical Engineering with a minor in Robotics at Georgia Tech in 2024, during which she was a member of Wish for Wash for four years and the co-VP of Recruitment for the team. Vibha's interests are in developing intuitive and adaptive wearable robots for clinical populations and hopes to go into academia

Mallika Digavalli is a fourth-year chemical engineering and material science minor at Georgia Tech, Her interests are centered around sustainability, and she hopes to eventually obtain a Ph.D. working in the field of developing biodegradable, biobased polymers that can serve as plastic alternatives. The following commentary served as a preliminary study to analyze the existing market, and this research is now focused on developing biobased, single-use menstrual products.

Arina Kucinskaja is a fourth-year biomedical engineering major with concentrations in chemistry and material science. She has served both as the co-Director and as the Menstrual Health Team Lead as a part of the Georgia Tech Wish for Wash chapter. Her interests are investigating the integration between chemistry, material science, and biomimetic designs with the hope to work on miniature therapeutics and/or surgical devices. She is currently in the process of pursuing a Master's Degree in Biomedical Engineering with potential transition into the PhD program

Abstract

Globally, 45 billion menstrual products are used annually. To reduce the environmental impact of waste accumulating from non-degradable plastics in discarded menstrual products, improving their sustainability is paramount. Sustainable menstrual products have been developed with strong consumer bases, as demonstrated by certain European Union (EU) and United States (US)-based companies. This paper focuses on the extent to which the products of different brands and companies in the EU and US compare in sustainability. The following companies: Freda, Corman S.p.A, Proctor & Gamble, Kimberly-Clark, Lil-Lets, and Honey Pot are analysed to evaluate the plastic usage in their menstrual products. These companies were selected through a rapid desk review and company-specific comparative analysis. This preliminary market research, limited by publicly available data, indicates a tendency for EU and small US companies to emphasise sustainability of their products more than large US companies. From our review, companies based in the EU have shown that it is possible to profitably manufacture sustainable menstrual products at scale. To improve global sustainability in the menstrual health sphere, we believe that US companies must follow suit.

Keywords: menstrual health, sustainability, United States, European Union, plastics

[1, 2, 3, 6]Biomedical Engineering, Georgia Institute of Technology, Atlanta, United States; [4, 5]Chemical & Biomolecular Engineering, Georgia Institute of Technology, Atlanta, United States; [7]Manufacturing Energy Efficiency Research & Analysis (MEERA) Group, Oakridge National Laboratory, United States; [8]Wish for Wash, United States

Introduction

Menstrual health, also known as feminine hygiene products, represents a large portion of the commercial market for global water, sanitation, and hygiene (WASH). The global feminine hygiene product market is projected to grow from $38.18 billion in 2021 to $54.52 billion in 2028 (Fortune Business Insights 2023). Due to this anticipated growth, sustainability has become an increased concern, with stakeholders pushing brands to adhere to Environmental, Social, and Governance (ESG) criteria in order to combat accumulation of waste (Investopedia 2023). There are two types of menstrual products available in the market: traditional single-use disposable products like tampons and pads and reusable products such as the menstrual cup and disc. The overwhelming majority of US consumers use single use tampons and pads as their primary menstrual products (Harvard T.H. Chan School of Public Health 2023). This preference carries a significant environmental burden. On average, menstruators have 451 periods in their lifetime, averaging to the use of 15000 tampons or 10000 pads (Ahuja and Singh 2022; Chavez-MacGregor et al. 2008). Consequently, even with the rise of new sustainable products, the dominance of traditional tampons and pads represents a critical environmental challenge, as these single-use items contribute to non-biodegradable landfill waste. Understanding the sustainability practices in dominant global consumer markets, such as the EU and the US, where consumer trends and product details are well reported on can better illuminate larger trends in the menstrual health field. This commentary investigates differences in sustainability between menstrual products of select EU and US companies by comparing their sustainability efforts and transparency.

Methods

A rapid desk review, a process in which companies and products were found through key word searches, was performed in 2021; from this, a company-specific comparative analysis is conducted in this commentary. This paper focuses on EU and US-based for-profit companies that manufacture pads and tampons, as they are the most common form of single-use menstrual health products. To identify the companies for the analysis, we researched over 20 companies, narrowing down the number of companies based on transparency. We define transparency as the degree to which companies provide verifiable, publicly accessible data regarding material compositions, supply chain ethics, and measurable environmental impacts to back their sustainability claims.

The selected companies are categorised into large and small companies. To standardize between varying international metrics of company size, we chose to define large companies as companies part of a

multinational corporation and conversely define small companies as companies based in a single country. It should be noted that the data analysed in this paper is based on publicly available information, including company websites. For the US, the second-largest market for feminine hygiene products worldwide (Statista Research Department 2022), two US-based multinational companies - P&G, with brands such as Always, Always Discreet, Just, Tampax, and This is L, and Kimberly-Clark, with brand Kotex - were selected due to their dominance in the global menstrual health market(Research and Markets, 2017). These two large companies were then compared with a smaller company in the US for a more holistic view of the market. Note that we identified one small company, Honey Pot, in comparison due to the lack of data available by other small companies. The three US companies were then compared to three similarly defined companies (two large, one small) in the EU. To identify the EU companies, we identified two of the largest European-based companies with menstrual product brands at the time, Premier FMCG, Corman S.p.A., and one small company, Freda (European Commission 2003). While Premier FMCG and Freda are based in the UK, they are categorized in this paper as EU entities because they maintain compliance with EU-based sustainability and ingredient transparency policies.

The basis of comparison across these six entities consists of their shared function of production of single-use menstrual hygiene products. To ensure equitable assessment of different companies, the analysis of corporations with multiple brands focuses specifically on their most eco-friendly brands (such as P&G's This is L.). This allows for a direct comparison between the best-in-class sustainable offerings of large companies and the models of smaller companies.

In each case study, the products with the most and least plastic content are highlighted. Specifically, analysis is centred around polyethylene, polypropylene, and polyester, as these are the most common unsustainable materials found in these products. Polyethylene and polypropylene are the typical manufactured plastics in products and can take several years to decompose, leaking toxic substances into the soil (Ujeh 2018; Jemec Kokalj et al. 2022). Polyester , an oil-based plastic that can take hundreds of years to decompose, is the leading contributor to microplastics in the ocean (Shah 2021). Additionally, we evaluated the use of bioplastics as an alternative ingredient for more sustainable products. While bioplastics have been used as a method to improve sustainability, it is important to note that not all bioplastics are considered "eco-friendly" (Ali et al. 2023). However, they can serve as a guide and intermediate towards more sustainable products. To be identified as a bioplastic, a material has to be composed of materials derived from living organisms and made up of organic compounds (Redwing, n.d.). Generally, the degradation rate of bioplastics tends to be faster than that of petroleum-based plastics (Rahman et. al. 2021). In this commentary, we define bioplastics as any plastics made from bio-based materials instead of petroleum (Gibbens 2021).

Finally, each company was compared based on the same 5 criteria. First, we assessed the absence of traditional plastics (polyethylene/polypropylene/polyester) within the product cores, individual packaging, and tampon applicators. Secondly, we evaluated the use of bioplastics and the level of ingredient transparency.

Results

Overall, the results show that the EU case studies met more criteria than the US case studies (Table 1). Across all companies studied, no plastics were found in the absorbent cores of the pads. However, results for individual product packaging varied. Only Freda offered plastic-free packaging for both tampons and pads. Lil-Lets achieved this solely for their tampons. A significant regional divide was observed in applicator materials and the use of alternatives. All EU companies reported plastic-free tampon applicators, while the US companies continue to utilize these plastics in their applicator designs. Furthermore, while two EU companies report the use of bioplastics in pads, the US companies demonstrated higher levels of ingredient transparency. Both P&G and Kimberly-Clark (US) provided the specific purpose for each ingredient listed, but only Lil-Lets (EU) had this information among the European companies (Table 1).

Company	Product	No plastic in core	No plastic in individual product packaging	No plastic in tampon applicator	Bioplastic usage	Purpose of each ingredient listed
Freda*	Pads	✓	✓		✓	
	Tampons		✓	✓	✓	
Premier FMCG- Lil-Lets	Pads	✓			✓	✓
	Tampons		✓	✓	✓	✓
Cormon S.p.A - Organyc	Pads	✓			✓	
	Tampons			✓	✓	
***Honey* Pot**	Pads	✓			✓	
	Tampons					
Proctor and Gamble - *This is L.*	Pads	✓			✓	✓
	Tampons					✓

Table 1 cont.

Company	Product	No plastic in core	No plastic in individual product packaging	No plastic in tampon applicator	Bioplastic usage	Purpose of each ingredient listed
Kimberly-Clark - *U by Kotex*	Pads	✓			✓	✓
	Tampons					✓

Table 1. Table representing different evaluating criteria in small (*) and large companies in both the European Union (underlined) and the United States (italics) when applicable for each of the two categories of products: pads and tampons. Dark cells indicate a category that is not applicable to the product, empty cells indicate an unmet criterion, and a check mark indicates a met criterion.

Freda, a small UK-based company (Freda n.d.) focuses on natural products that contain organic and renewable materials. Freda's stated mission considers the environmental impact of its products as it emphasises that 75-80% of its products are made from renewable materials. Regarding the criteria identified for this study, Freda's pads have a wood pulp core and a bioplastic packaging; their tampons use bioplastic packaging and cardboard applicators. While they use wood pulp and bioplastics, they provide no details on proportions, specific materials, or their roles.

Lil-Lets, a large UK-based company acquired by multinational Premier FMCG (Pty) Ltd, promotes its participation in the EU Single Plastics Directive. While Premier FMCG has many brands, Lil-lets is its only feminine hygiene brand. It is a top global market contributor (Research and Markets 2017), prioritising sustainability with the goal of "protecting the planet with every period" (Lil-lets n.d.). Aligned with the EU Single Plastics Directive, Lil-Lets packaging details plastic composition and proper disposal methods. Related to the criteria evaluated in this study, their pads use a cellulose core, non-applicator tampons are cellophane wrapped, and they offer tampons with cardboard or no applicators. As for plastic alternatives, they use bamboo, cellulose, and PLA. Finally, they show high ingredient transparency by describing each

material and its purpose, and specifying the amount of plastic per individual product.

Organyc, based in Milan, Italy, is a large company under the multinational Corman, (Organyc n.d.). While Corman has multiple brands focusing on feminine hygiene, Corman is the only brand that does both pads and tampons. Its main goal is to provide clinically safe products for women. Related to our criteria, their pads use organic cotton core, they mention packaging materials as "100% biodegradable" but do not specify any specifics, they offer cardboard and bioplastic tampon applicators, and have corn-starch-based bioplastics used for protective sheets in pads. While Organyc claims to have both products and packaging that are "100% Biodegradable," they provide no traceable proof or quantification to verify this claim.

As for the U.S. companies, **Honey Pot** is a small company founded in 2014 focusing on plant and naturally derived formulas (The Honey Pot n.d.). Honey Pot lists a description of its ingredients, and is going further, campaigning for a legislative movement requiring comprehensive ingredient lists and ingredient transparency (Scranton et al. 2022). As of 2024, Honey Pot has been acquired by Compass Diversified (CODI), a public trading holding company. As this paper was in review and the company acquisition has not been completed at this time, the authors still consider Honey Pot as a small US company. The acquisition is not predicted to alter any product formulas as of now. Related to our criteria: their pads use a wood pulp polymer-based core, packaging uses plastics, their tampons use BPA-free plastic applicators, and they do not specify the purpose and usage of product ingredients.

Procter & Gamble (P&G) is a large US-based company specialising in a wide range of hygiene products (Proctor and Gamble 2023). This multinational corporation owns five menstrual product companies: Always, Always Discreet, Just, Tampax, and This is L. Within those, Always and Tampax are the number one menstrual product sellers in the US (Petruzzi 2022). P&G's most publicly promoted sustainable brand is This is L. In terms of criteria, the following information for the selected criteria was summarised based on This is L. due to their promoted sustainability: their pads use cellulose for the core, tampon have plastic packaging, tampon applicators are BPA-free plastic, and they provide ingredient transparency showing each material and their purpose.

Finally, **Kimberly-Clark** Corporation is a large US company with the goal of making "lives better ...with the smallest environmental footprint." (Kimberly-Clark n.d.) Most feminine care products fall under the brand Kotex. In the US, this is mainly through U by Kotex. Kimberly-Clark differentiates brands based on location, not based on types of products. They partially meet our sustainability criteria: their pads have wood pulp cores, tampons use plastic applicators, all items are individually wrapped in plastic, and ingredients are clearly disclosed.

Limitations

It should be noted that the information in this paper is limited by publicly available data, and as a result, these specific case studies were chosen. As only six companies were discussed, there is ample opportunity in the future for a deeper analysis into regional sustainability trends. Ingredient analysis comes from review of company websites which may not be fully transparent. Finally, definitions of company size differ between the two regions. Therefore, companies were compared relative to these criteria rather than overall global market impact. This commentary is a starting point for a more sustainable perspective and provides a base for future studies to dive deeper into sustainability issues of menstrual health products.

Discussion

Overall, these case studies demonstrated that EU companies met more of the identified sustainability criteria than those of the US. For example, Honey Pot, a US-based small company, has put its focus on naturally derived, plant-based ingredients that have less detrimental impacts on the environment. The usage of bio-plastics made from sugarcane bolsters this claim, but there is still room for improvement due to their polyethylene applicators. These applicators have BPA-free, low-density polyethylene, but Honey Pot's sustainability efforts would benefit from eliminating polyethylene entirely. By comparison, Freda, an EU-based small company, focuses on renewable materials to reduce the environmental impact of their products as seen in their Sustainability Statement. Freda period eco pads have 75 - 80% renewable materials and all

their pads use wood pulp as the absorbing material. Additionally, their tampons come in cardboard applicators, which is a stark difference in sustainability measures. When it comes to transparency, both of the small companies provide less material disclosure than most of the large companies.

Looking at large companies, none of Lil-let's tampons have plastic applicators, but the majority of P&G's Tampax products do. The Lil-Lets tampon with the highest plastic content still contains less plastic than most Tampax products. P&G has demonstrated an ability to improve sustainability through cotton-based and recyclable menstrual products with its brand This is L. However, their sustainable impact is heavily limited without transitioning this to their other brands and with the consistent use of plastic in their products. Like P&G's This is L., Organyc tampons and pads are made with cotton. They also integrate recycled cotton into their tampons, which have either a cardboard or plant-based bioplastic applicator. Additionally, their pads have a biodegradable protective sheet. In contrast, P&G lacks this integration of sustainable materials. Additionally, Kimberley-Clark's products are most similar to P&G's in terms of sustainability: their pad's core contains wood pulp, but both the packagings and tampon applicators are plastic-based. These companies, all with considerable impact on the menstrual health product market, have striking differences in their commitment to sustainability.

When investigating why these companies produce such different results in sustainability, one must look no further than the regulations. The UNI EN 13432 standard, a compostability standard for packaging, exemplifies one of the key reasons why EU-period companies can create biodegradable products that are better for the environment ("Compostability – BIOPAP" n.d.). These standards require at least 90% disintegration after 12 weeks, 90% biodegradation in six months, and tests on ecotoxicity and heavy metal content. In comparison, US standards are set by the American Society for Testing and Materials; for materials to be considered biodegradable, ASTM D6400 and D64848 state that 60 percent of the plastic must be broken down within 180 days in a commercial composting facility. This essentially negates the idea of plastic products being biodegradable and compostable, as most people do not have access to a commercial composting facility (Biosphere n.d.). Furthermore, there are key differences in the policy approach for sustainability in the EU and US.

In 2019, the EU generated 53 million metric tons of plastic waste and this number continues to increase (Alves, 2023). A new wave of sustainable legislation aims to reduce this, specifically via the EU Directive on Single-Use Plastics. Enacted on July 2nd, 2021, the directive outlines general guidelines for plastic use in a variety of industries, from food packaging to sanitary products (Gesley 2021). More specific menstrual health guidelines arise from individual countries and companies incorporating reusable and biodegradable materials within their products (Official Journal of the European Union 2019). While the EU has begun implementing fundamental legislation to regulate plastic use and production, the same cannot be said for the US. While the large US companies we examined demonstrate strong material transparency, companies have relative freedom in the production of menstrual health products due to a current lack of comprehensive and regulated national sustainability standards. Congress has been trying to pass S.984 - Break Free From Plastic Pollution Act of 2021 to gradually ban some single-use products (Congress.gov 2022). Even with introducing this Act, the US is behind in plastic management and regulation when compared to the EU. The US produced approximately 35.7 million tons of plastic waste in 2018 (Environmental Protection Agency 2023). Menstruators in the US discard twelve billion pads and seven billion tampons every year (Stanford Magazine 2017). Due to the lack of regulation from the US government, companies alone decide their involvement in sustainable production.

One argument in favor of these US companies centers around cost as a barrier to producing more sustainable products, as evidenced by companies designating these products as 'premium'. However, these companies are not restricted in expenses. P&G and Kimberly Clark have invested millions around the world for period poverty initiatives. Though this is empowering millions of menstruators around the world, sustainability is too pressing an issue for these companies to restrict their philanthropic efforts to just one avenue.

Conclusion

In both the EU and the US, companies, especially large ones, have both the resources and the technological capabilities to produce sustainable menstrual products. Therefore, the question is no longer if they can, but why they will not. Companies in the EU have proven that it is possible to produce sustainable products whilst making a profit. While the large US companies we examined demonstrate strong material transparency, which should be maintained moving forward, the amount of plastic in the US menstrual products is notably more than those produced by the US. Consumers are holding companies accountable for reflecting this priority in their products (Frey et al. 2023). Sustainability can no longer be ignored. Companies can join the global sustainability movement by prioritising more sustainable ingredients and moving away from single-use plastics in their products. It is the responsibility of brands to work towards a more sustainable future, one that many have already taken on. The US can achieve this level of sustainability if they prioritise a platform of sustainability – a commitment that is paramount in today's world.

Acknowledgements

The Authors of this paper would like to thank Aline Banboukian, PhD., and Jasmine Burton, MPH, MBA, for serving as our reviewers and editors during the process. We would also like to thank Jasmine Burton, Prerana Kumar, and the whole of Wish for Wash for serving as great resources, mentors, and support systems during our work on this commentary.

Declaring Conflict of Interest

We have no known conflict of interest to disclose. This paper does not promote nor receive monetary benefits for choosing one company over another. Nor should the use of these companies be seen as a promotion of one company's products over the others. This paper serves as an undergraduate research study and not for commercial evaluations. These case studies are solely used to describe the general situation of each entity given the limited public information.

Access references for this paper using the QR code below

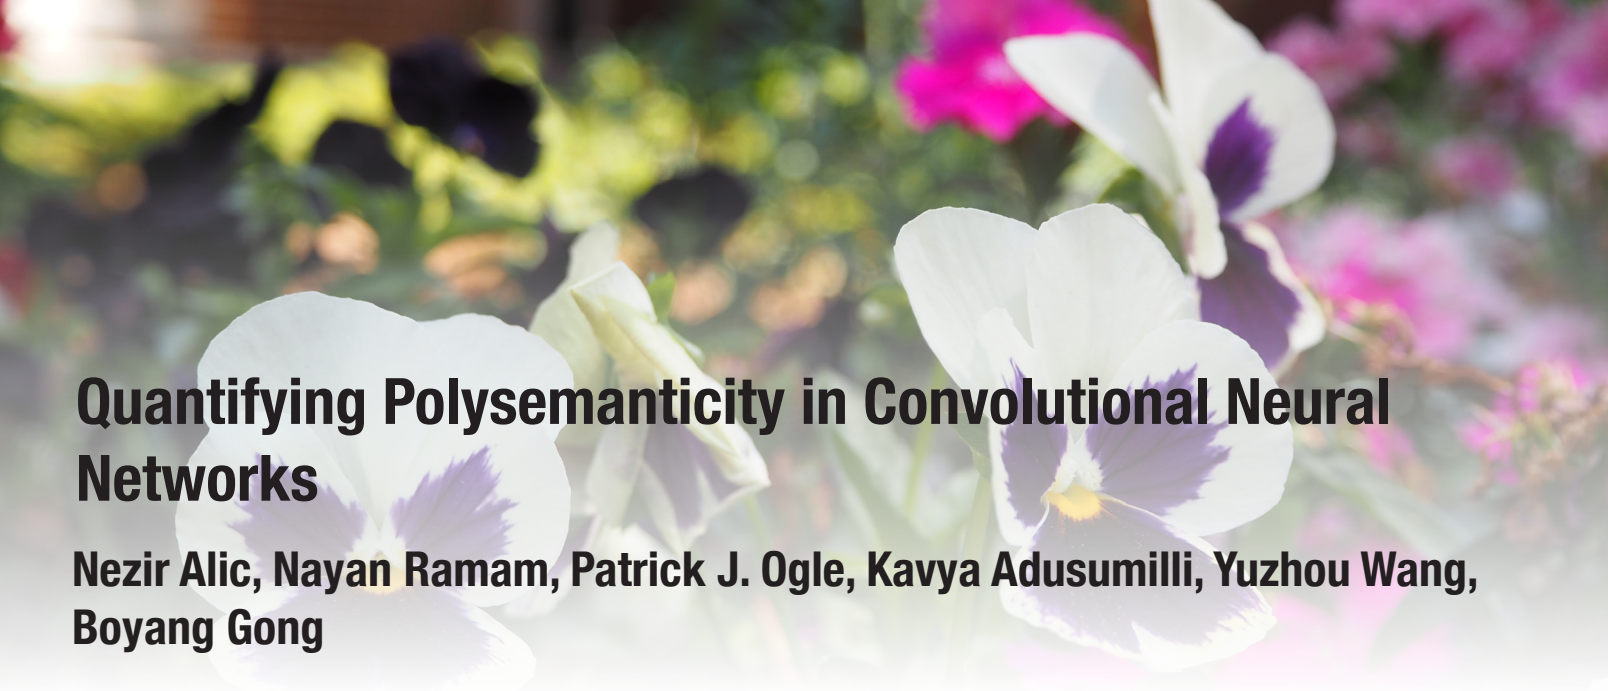

Quantifying Polysemanticity in Convolutional Neural Networks

Nezir Alic, Nayan Ramam, Patrick J. Ogle, Kavya Adusumilli, Yuzhou Wang, Boyang Gong

Nezir Alic is a 4th year PhD student in physics, currently running GPU-accelerated simulations of black hole accretion disks in order to better understand the dynamics of super-Eddington accretion. He has a wide range of interests and plans to transition to a role in industry after his PhD. Nayan Ramam is a second-year Electrical Engineering major with interests in analog design and machine learning. Nayan works in Dr. Jennifer Hasler's Integrated Computational Electronics Lab, is on the Analog Mixed-Signal and Digital Design sub-teams of Silicon Jackets, and is a project lead in the Interactive Music Group VIP. Upon graduation, he aspires to work in analog VLSI – neuromorphic circuitry articularly captivates him. Patrick Ogle is a second-year Physics and Mathematics double major. He is currently a part of the Grigoriev Research Group in the Georgia Tech Department of Physics.Though not having a concrete grasp as to what his professional goals are, something in academia or academia-adjacent will probably await him. Kavya Adusumilli is a third-year Computer Engineering major interested in machine learning, probabilistic modeling, and quantitative finance. Her work focuses on developing and evaluating algorithmic trading systems, and she plans to pursue a career at the intersection of AI and financial markets.Yuzhou Wang is a Ph.D. candidate in the Algorithms, Combinatorics, and Optimization (ACO) program, with research interests in the interface of probabilistic combinatorics, statistical physics and algorithms. His work explores random structures,—such as random (hyper)graphs and spin glasses, along with related questions in inference, concentration, and algorithmic limits. Boyang Gong is a second-year Industrial Engineering major with a concentration in Data Science. He is currently a researcher at AI4Science @ Georgia Tech on the VizFold project, where he develops interactive visualizations to improve the interpretability of protein folding mechanisms within OpenFold architectures. His future research directions involve refining metrics for architectures to enhance predictive modeling in other area.

Abstract

Discerning between polysemanticity, the phenomenon whereby individual neurons respond to multiple unrelated features, and monosemanticity, the phenomenon whereby individual neurons respond to only one feature, poses a challenge for neural network interpretability. While previous research has relied on subjective visual inspection to assess whether neurons are monosemantic or polysemantic, no systematic quantitative framework has been established. We present a novel automated approach for measuring polysemanticity in convolutional neural networks (CNNs) through the analysis of feature visualizations. We develop a scoring formula that combines cluster count and angular diversity measures to produce normalized polysemanticity scores ranging from 0 (perfectly monosemantic) to1 (maximally polysemantic). Validation on a modified LeNet-5 architecture trained on the CIFAR-10 dataset demonstrates that our feature visualizations consistently achieve higher activation scores than top dataset examples, confirming effective identification of optimal activation patterns. Our framework provides an objective, automated method for quantifying neuron polysemanticity, enabling systematic studies of how training procedures and architectural choices influence interpretability. This work establishes a foundation for developing more transparent neural networks and advancing mechanistic understanding of learned representations.

1 Introduction

Neural networks have revolutionized machine learning, achieving remarkable performance across diverse tasks. These particular tasks, such as diagnosing illness from numerous medical data and driving cars autonomously, have proven to be incredibly effective implementations of neural networks. However, the increasing use of neural networks in day-to-day life in evermore critical applications has redoubled efforts in understanding the inner workings of neural networks, a subject fogged in light of the "black box" view this technology (Y. Zhang et al. 2021). One particularly interesting aspect of neural network interpretability (the understanding of specific neuron behavior) concerns the degree of polysemanticity (when one neuron responds to several unrelated features) in individual neurons (Elhage et al. 2022).

Previous research has explored neuron activations through visualization techniques (Olah, Mordvintsev, and Schubert 2017) and activation analysis (Scherlis et al. 2025), such as ascribing to interpretable features directions present in neural networks (i.e. embeddings) (Elhage et al. 2022). However, determining the degree of polysemanticity in a neuron has remained largely subjective and qualitative. This lack of objective measurement hampers the systematic study of how training choices and architectural decisions affect semantic representations within neural networks. By providing a quantitative measure of polysemanticity, we allow for automated and quick insight into neural networks, easing the process of developing interpretable networks.

We adopt the convention that features are the fundamental unit of neural networks, the smallest unit of human-interpretable meaning carried by a particular neuron (Olah et al. 2020). By this "fundamental unit" we mean that a feature is a property of the input that a sufficiently large neural network will reliably dedicate a neuron to. With this viewpoint adopted, we take it upon ourselves later on in this paper to rigorously study these features.

The driving principle behind polysemanticity is superposition, which is the phenomenon that occurs when a neural network aims to encode more features than there are neurons (Scherlis et al. 2025). In essence, the comparatively small neural network is a "projection"of a much larger network, accomplished by encoding linear combinations of several features into single neurons. To this end, the aforementioned features are represented by directions in a privileged basis, as an incentive to align with basis directions.

In this paper, we introduce a novel metric for quantifying the semanticity of neurons in convolutional neural networks. By combining feature visualization techniques with image similarity measures, we develop

a method to score neurons based on the diversity of patterns that activate them. This approach provides an automated, objective way to measure semantic properties without requiring manual inspection.

This work contributes to the understanding of neural network interpretability by providing a quantitative framework for analyzing the semantic properties of neurons and offers practical insights for designing more interpretable models.

2 Methodology

Our motivating intuition behind characterizing a neural network's semanticity stems from the fact that polysemantic (multi-responsive) neurons should ideally have feature visualizations that are visually diverse, i.e. having multiple discernible features. To this end, we created an automated routine of clustering neurons in embedding space, analyzing this cluster distribution, with the aim of determining quantitatively the activation diversity of a given set of neurons.

Each neuron in a neural network requires different input characteristics to activate. That is, different neurons "respond" to different aspects of the input. It is possible to create images that maximally activate a neuron using gradient ascent, a process called feature visualization (Olah, Mordvintsev, and Schubert 2017; Simonyan, Vedaldi, and Zisserman 2014). If the range of characteristics that activate a neuron are relatively narrow, then one can expect the feature visualizations to reflect this by being similar to each other. This case corresponds to a high degree of monosemanticity (or equivalently, a low degree of polysemanticity). If, on the other hand, there exists a wide array of features that activate a neuron, then the feature visualizations can be expected to differ more greatly from each other. This case corresponds to a high degree of polysemanticity. It follows that it should be possible to gain some information about the polysemanticity of the neurons in a network by examining the extent of the variation in feature visualizations. Figure 1 outlines our implementation:

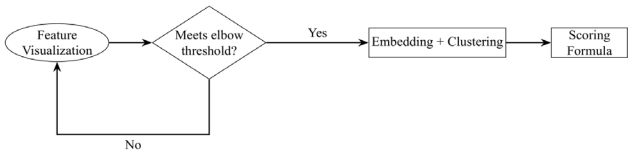

Figure 1. Block Diagram of the Neuron Polysemanticity Scoring Pipeline

2.1 Feature Visualization

Feature visualization is a technique used to understand what individual neurons in a neural network respond to by identifying input patterns that strongly activate them. A common method involves starting with a random image and applying gradient ascent to adjust the pixel values in order to increase the activation of a specific neuron. This process works to reveal the features that the neuron has become sensitive to during training. Regardless of the approach, the goal is to make the model's internal behavior more interpretable by showing the kinds of structures, textures, or concepts each neuron is tuned to detect.

2.2 Elbow Detection

In order to capture the entire activation space of a neuron, it is critical to create enough feature visualizations. We do so by first generating two images and finding the similarity between them using the Learned Perceptual Image Patch Similarity (LPIPS) metric (R. Zhang et al. 2018). We then generate an additional image, and find the absolute change in the average image similarity. If the magnitude of this change falls below a threshold (the elbow), we stop generating new images. If the magnitude does not fall below the threshold, we generate another image, yielding a new change in average similarity. This process is repeated until the threshold is met. We empirically determine that a threshold of 0.02 (i.e. 2% change) achieves a balance between comprehensive coverage of the activation space and computational efficiency. Lowering the threshold beyond 0.02 produced diminishing returns with near-duplicate visualizations. Higher thresholds risk undersampling diverse activation patterns. This threshold yielded 4-10 visualizations per neuron on average for our LeNet-5 architecture.

2.3 Lpips - Learned Perpetual Image Patch Similarity

To quantify similarity between images, we use the LPIPS metric, a framework based on human visual perception that aligns more closely with how people assess visual similarity (Ghazanfari et al. 2023). Unlike pixel-wise metrics, LPIPS compares images in deep feature space, capturing perceptual differences rather than low-level pixel variations. To compute LPIPS, we pass each image pair through a pretrained convolutional neural network. LPIPS extracts feature activations from several layers of the network, normalizes them channel-wise, and compares them using a weighted L2 distance. The weights are designed to accurately reflect human perceptual judgments. We compute LPIPS using a Python implementation by Zhang et al. (2018).

2.4 Embedding and Clustering

To prepare our feature visualizations for clustering, we use a ResNet-18 architecture pretrained on ImageNet for embedding. By passing through the feature visualizations and extracting before the final fc layer, we obtain a 512-dimension embedding for each visualization. We chose ResNet-18 for its low memory and compute footprint. Embedding resolution (512-D as opposed to other models with 1024 or 2048-D embeddings) was not a huge concern given the size of our toy model and the relative simplicity of the CIFAR-10 dataset. Upon expansion of our methodology to larger networks, we may consider switching to another CNN architecture with higher capacity.

An important feature of the embeddings we obtain is that each dimension corresponds to the same abstract feature across each visualization in a given neuron. We take advantage of this fact in our scoring formula. Another added benefit of embedding our visualizations is that it increased computational efficiency, as the sco ing process can work with 512-D vectors as opposed to the feature visualizations, which are 1024-D tensors.

We then perform KMeans clustering to sort the embeddings into groups. We determine the optimal number of clusters using the Bayesian Information Criterion (BIC). The BIC term functions by minimizing the sum of the inertia (a loss term, the sum of squared distances from cluster centroids) from KMeans clustering and a penalty term proportional to the number of clusters. In KMeans clustering, it is certain that the inertia will decrease as the number of clusters is increased. To pick the number of clusters that balances inertia and cluster count, the BIC employs the aforementioned penalty term to punish an increase in cluster count. In practice, this process "groups" similar images into clusters, with each cluster then representing a distinct feature for a neuron. Clustering accounts for the case where many near identical images are created, all of which represent the same feature in the neuron's activation space but would skew the polysemanticity score upward due to imperfections and diversity in gradient ascent. This allows us to overproduce feature visualizations by lowering the elbow threshold. This is desirable since it ensures more opportunities to generate diverse feature visualizations and improves coverage of the neuron's activation space.

2.5 Scoring Formula

From the clustering step, we collect two critical pieces of data: the number of clusters K and the centers of each cluster $\{C_i\}_{i=1}^{K}$ .Given these, we propose a six-step process to determine a value in [0, 1] that accurately quantifies the polysemanticity of a neuron.

2.5.1 Norms And Unit Vectors

Define m_i as the magnitude of the ith cluster center vector, and ni as the unit vector for the ith cluster center vector:

$$m_i = ||c_i||_2, n_i = c_i/m_i$$

2.5.2 Center Vector Angles

Define $\theta_{ij} \in [0, 1]$, the normalized angle between n_i and n_j, as

$$\theta_{ij} = \frac{1}{\pi} \arccos(n_i \cdot n_j)$$

It is not difficult to see that θ_{ij} measures directional dissimilarity between two clusters.

2.5.3 Magnitude Weight Term

Define $b_{ij} \in [0, 1]$ as

$$b_{ij} = \frac{2min(m_i, m_j)}{m_i + m_j}$$

In this way we can downweight pairs where one cluster is much weaker, as a second cluster

being weaker indicates the second feature representing a smaller part of the neuron's activation space.

2.5.4 Angulur Diversity

Define the angular diversity D as the mean of $\theta_{ij} b_{ij}$ over all distinct pairs of clusters:

$$D = \frac{1}{\binom{K}{2}} \sum_{i<j} \theta_{ij} b_{ij}$$

2.5.5 Cluster Count

Normalize the number of clusters K based on the number of feature visualizations I:

$$C = \frac{K-1}{I-1}$$

Notice that C = 0 if K = 1, meaning all visualizations are in one cluster; and C = 1 if K = I, meaning every visualization is "unique."

2.5.6 Final Score

The final polysemanticity score S for that neuron is defined as

$$S = 1 - (1-C)(1-D) = C + D - CD$$

This formula behaves like the probability of at least one event occurring in a union—much like

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

where high S (approaching 1) emerges only when both multiplicity and diversity are substantial, signaling a neuron responding to many unrelated features; moderate S if one factor dominates (e.g., many clusters but all similar, or few but very diverse); and low S (near 0) when both are minimal, suggesting monosemantic behavior with a single dominant feature. This interaction ensures S penalizes noise or superficial variety while amplifying genuine polysemantic complexity, all normalized to [0, 1] for interpretability.

Table 1 depicts the score range that our formula permits accompanied with an approximate interpretation of each.

Score Range	Interpretation
0.0 – 0.3	Highly monosemantic
0.3 – 0.5	Moderately monosemantic
0.5 – 0.7	Moderately polysemantic
0.7 – 1.0	Highly polysemantic

Table 1: Approximate Polysemanticity Score Interpretation.

3 Results and Limitations

To corroborate the validity of this metric, we verify the ability of the feature visualizations to capture the top activating feature(s). We feed our feature visualizations through the neural network, along with a batch of CIFAR-10 images, gathering the activation scores for all images on the neuron being scored. By simply comparing the activation scores for the feature visualizations to those of the CIFAR-10 test images, we can determine whether the gradient ascent process was effective in finding the top activating feature(s). We observed that our feature visualizations consistently yielded activation scores higher than those of the CIFAR-10 test images, indicating consistent deduction of the main activating features.

We analyzed 394 neurons (every neuron in our network), observing an overall average polysemanticity score of 0.350, indicating general monosemanticity in the network. This aligned with our qualitative inspection of the output. We did not observe strong evidence of a variation in polysemanticity by layer, but a robust conclusion would require further testing. In Figure 2, we present the distribution of polysemanticity scores by layer, as well as an example of a polysemantic and monosemantic neuron. Moreover, Figure 3 depicts

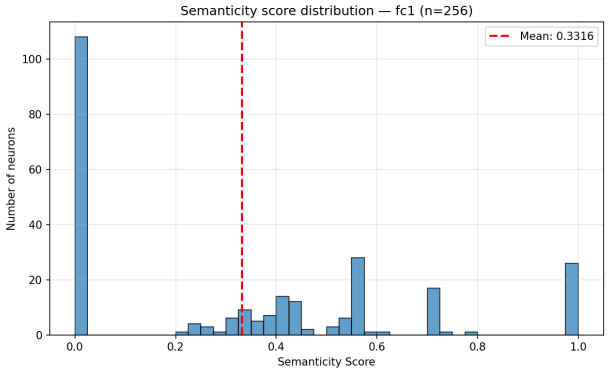

(a) fc1

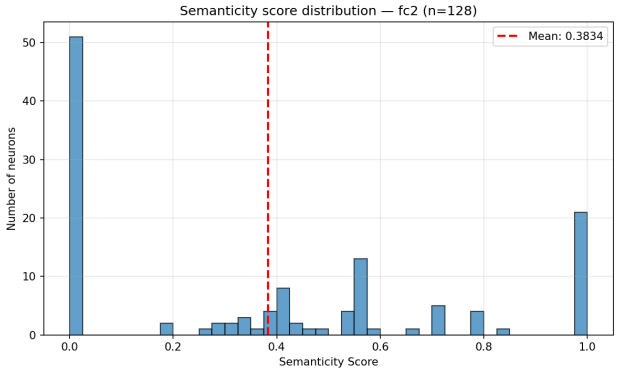

(b) fc2

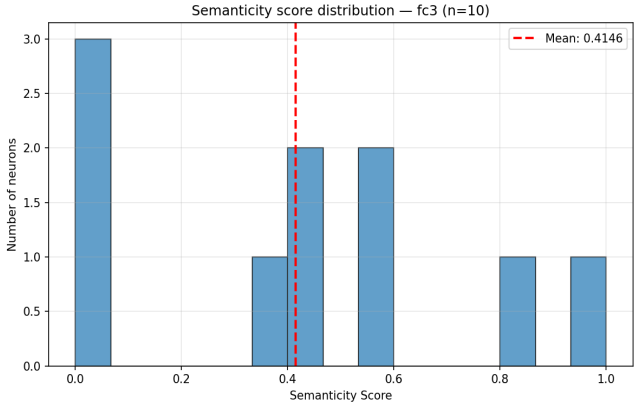

(c) fc3

Figure 2: Distribution of semanticity (polysemanticity) scores across neurons for fully connected layers fc1, fc2, and fc3. Each histogram shows the number of neurons at each score; the red dashed line indicates the layer mean.

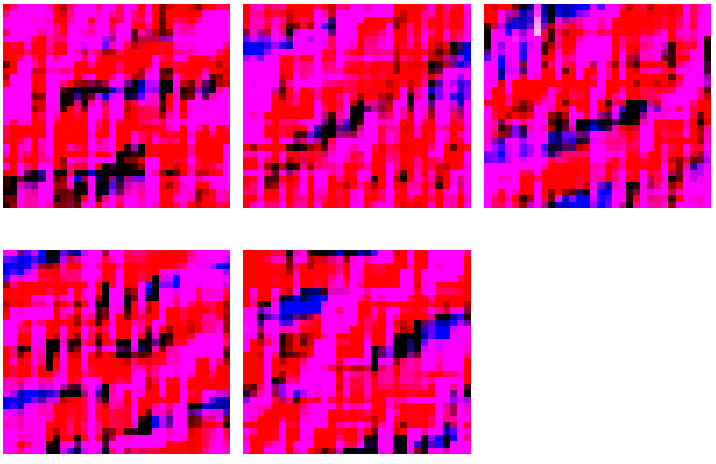

(a) Feature visualizations for a monosemantic neuron (polysemanticity score of 0.0). Note that they are all nearly identical.

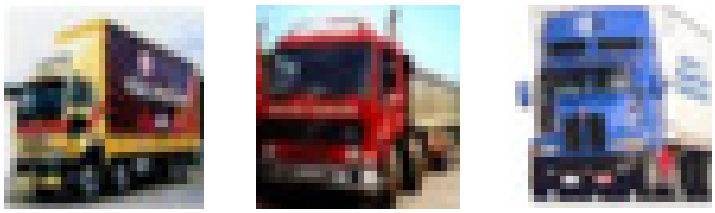

(b) Images from CIFAR-10 test batch with highest activation for the same neuron

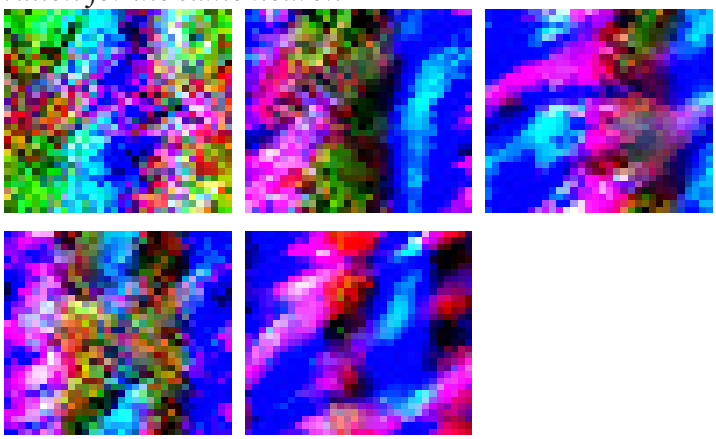

(c) Feature visualizations for a polysemantic neuron (polysemanticity score of 1.0). Note that the image variance is qualitatively much higher than the monosemantic neuron.



(d) Three of the highest activating CIFAR-10 images for the neuron from Fig. 3c. This neuron saw high activation from dogs, birds, automobiles, and more, aligning with the high polysemanticity score.

Figure 3: Comparison of monosemantic and polysemantic neurons.

selected monosemantic and polysemantic neurons.

4 Discussion

This work presents a novel framework for quantifying polysemanticity in convolutional neural networks through automated analysis of feature visualizations. Our approach addresses a fundamental challenge in neural network interpretability: moving beyond sub jective, qualitative assessments of neuron behavior to provide more objetive, quantitative measures of semantic properties.

Our key contributions include: (1) a comprehensive

pipeline that combines feature visualization, embedding-based clustering, and geometric analysis to measure polysemanticity; (2) a normalized scoring formula that accounts for both cluster multiplicity and angular diversity in embedding space; and (3) systematic validation methods that verify the quality and coverage of generated visualizations.

The scoring formula $S = C + D - CD$ effectively captures the intuition that polysemantic neurons exhibit both multiple distinct activation patterns and significant diversity between those patterns. By incorporating magnitude weighting and normalized angular measurements, our approach provides robust quantification across different network architectures and training conditions.

While our validation demonstrates successful identification of optimal activating patterns, several limitations remain. The approach is currently validated on a modified LeNet architecture trained on CIFAR-10, and extension to larger, more complex architectures requires further investigation. Additionally, the reliance on ResNet embeddings for clustering introduces potential biases that may not fully capture the semantic organization of other network types.

The main focus of our future work is to perform large scale testing to further validate the robustness of our metric. We also plan on applying our framework to study how different training configurations affect the emergence of polysemanticity during learning, investigating the relationship between semanticity scores and model performance, and extending the approach to transformer architectures and other neural network families. The development of automated, objective measures for neural network interpretability represents a crucial step toward building more transparent and trustworthy AI systems.

4.1 Future Works

We implement a custom Wasserstein GAN-CGAN (Arjovsky, Chintala, and Bottou 2017; Mirza and Osindero 2014) hybrid to enable the creation of conditioned datasets based upon the 10 labels that characterize CIFAR-10. A custom training pipeline was used to allow the insertion of custom 32x32 masks into the inference time of the WGAN-CGAN. This masking feature provides a qualitatively sound way of embedding controllable and consistent contours/shapes within the dataset, opening new ways of empirically gauging semanticity. Though not finalized in its robustness to be used in this project, such an endeavor to completed and applied to this project appears to be fruitful in further corroborating our devised semanticity formula.

Acknowledgements

We would like to thank the Math Modeling Student Research Group at Georgia Tech for creating the opportunity for the authors to meet and embark on this project. We also thank Zach Furman, who is pursuing doctoral studies in mechanistic interpretability at the University of Melbourne, for useful discussions and advice in planning this project.

Access references for this paper
Access references for this pa-



The Impact of the 2021 Euro-5 Fuel Standard on PM2.5 Levels in Karachi

Dan V. Galea[1], Jinho Kim[2], Marika Pescosolido[3]

Dan V. Galea, Jinho Kim, and Marika Pescosolido are students at the Georgia Institute of Technology. Dan V. Galea is a third-year undergraduate majoring in Electrical Engineering, completing a minor in Public Policy and interested in energy system and expansion of the power grid. Jinho Kim completed a degree in Biology, graduating in Fall 2025 and is now enrolled in a Master of Science program in Bioinformatics. Marika Pescosolido is a first-year student pursuing a degree in Physics with a minor in German Language, fascinated by astrophysics, particle physics, and quantum mechanics.

Abstract

This study investigates and quantifies the effects of the 2021 Euro-5 Low-Sulfur Fuel Standard specifically on PM2.5 concentrations, in addition to NO_x , SO_2, and CO. It uses a rigorous Quality Control (QC) process to select the most suitable hourly PM2.5 time series data for Karachi (2019-2025). The analysis then constructs a synthetic counterfactual using data from unexposed neighbouring cities (i.e. donors) to estimate the counterfactual PM2.5 trend in the absence of the January 2021 Euro-V policy intervention. Using Difference-in-Differences (DiD) testing on the STL-adjusted data (season), the policy is found to have a statistically significant positive effect, resulting in reduced PM2.5 levels. The credibility of this finding is supported by a placebo test, confirming the Euro-5 policy's impact on air quality in Karachi.

Keywords: Karachi · PM2.5 · Air pollution · Synthetic counterfactual
Field of research: Environmental Regulation · Data Science · Public Policy

1,2,3Georgia Institute of Technology, School of Public Policy, Atlanta 30332, United States

1 Air Pollution and Its Health Implications in Karachi: Background and Significance

According to "Dawn", the PM2.5 levels recorded in Karachi are approximately 20 times the WHO's (World Health Organization) annual guideline value, placing the port city of Pakistan among the top five most polluted cities in the world.[1] Residents are advised to avoid outdoor exercise and rather prioritise wearing masks and/or running air purifiers in working environments and households. Smog is set to intensify in the coming years if no action is taken due to "a combination of local emissions and pollutants carried by low-speed winds from India". Undoubtedly, air pollution poses a hazardous risk to public health, increasing the chances of reporting respiratory issues The WHO Representative, Dr Dapeng Luo, has alarmingly stated that 256,000 people are dying in Pakistan every year due to air pollution. As it will be shown in the following pages, smog levels witness a spike during winter, contributing significantly to the annual deaths.[2] This seasonal pattern finds its explanation in the lower temperatures in the Northern Hemisphere due to air stagnation, generated by "atmosphere trapping". The latter is caused by temperature inversions, where dense cold air holds pollutants near the ground surface under the warmer air, acting like a blanket trapping particulate matter. Furthermore, winter brings rising emissions from winter heating, which includes the burning of wood and biomass, as well as increasing vehicle emissions (for analogous reasons), and weaker winds (i.e. natural air cleanings), slowing down pollutants disposals.[3] Moreover, increasing humidity due to Pakistan's warm, moist air from the Arabian Sea and Western Disturbances contributes to creating ideal conditions for "trapping" fogs. The warm, sea air, when it encounters the cooler temperatures in Karachi, condenses into water droplets, raising relative humidity, which causes particles to stay suspended longer.[4] Additionally, high humidity makes PM2.5 particles grow larger by absorbing the air water droplets previously mentioned. This effect is known as hygroscopic growth and worsens human health by changing the aerodynamic properties of the inhaled particles: larger particles tend to deposit in the upper airways, rather than alveolar regions. This negatively emphasises asthmatic symptoms, allergies, and COPD (Chronic Obstructive Pulmonary Disease), which causes airway narrowing. This explains why penetrating the extrapulmonary regions (i.e. upper regions) can be even more damaging.[5]

1.1 Implementation of Euro-5 Vehicle Emission Standards

In order to reduce the hazardous AQI levels, the Pakistan government approved, by January 2021, the transition from Euro-2 to Euro-5 for both petrol and diesel. The new refining policy was incentivised by the duty-free import of machinery, as declared by Sumit Ritolia, refinery economics analyst at S&P Global Commodity Insights. He also added "By prioritising the upgrading of existing refineries[6] to manufacture Euro-V grade fuels, the government underscores its dedication to aligning with global environmental benchmarks and bolstering energy infrastructure".[7,8] Compared to the previous Euro-2, now much lower limits for carbon monoxide (CO), nitrogen oxides (NO_x), Hydrocarbons (HC), and particulate matter (PM) are required, especially for diesel vehicles. Thus, Euro-5 drastically cuts the maximum allowable sulfur content (i.e. Low-Sulfur Fuel Standard) in diesel and gasoline to 10 parts per million (ppm).[9] This reduces the emissions of soot and smog, and enables the advanced emission-control technologies and fuel-efficient designs.

EURO-5: Emission Standards

The Euro-5 standard refers to the allowed emission values for the containment of particulate matter, carbon monoxide, and nitrogen oxides for vehicles using gasoline and diesel engines. The emission standards mentioned are shown below:

Pollutant	Emission (g/km)
CO	1
HC	0.1
NOx	0.06
VOC	0.068
PM	0.005 (direct injections)

2 Hypothesis

Our hypothesis posits that, relative to a synthetic counterfactual, the adoption of Euro-V fuel standards in Pakistan is expected to lead to a measurable reduction in PM2.5 concentrations in Karachi. Existing studies examining similar policy shifts in other countries have demonstrated a clear association between improved fuel standards and decreases in PM2.5 levels.[10,11,12] However, during the course of the analysis, a critical implementation challenge was identified: Pakistani domestic refineries have not yet been upgraded to produce fuel that meets Euro-V standards. This regulatory and infrastructural shortcoming limits the supply of compliant fuel within

the country.[13] Consequently, an alternate hypothesis is proposed — inadequate regulatory enforcement of domestic refineries has constrained the expected reduction in PM2.5, rendering the policy less effective than intended. Further evaluation was undertaken to assess the impact of the slow adoption of the EuroV policy. Literature indicates that petroleum, which Pakistan primarily imports and which must therefore meet Euro-V standards upon entry, tends to generate higher carbon monoxide (CO) emissions per capita. In contrast, diesel, predominantly refined domestically, produces substantially greater quantities of nitrogen oxides (NO_x) and sulfur dioxide (SO_2 per capita.[14,15] By examining trajectories of NOx, SO_2, and CO emissions in Pakistan against those of a synthetic Pakistan constructed from the donor pool of PM2.5 comparison countries, it is possible to analyze whether the discrepancy between imported and domestically refined fuels has hindered the intended effects of Pakistan's Euro-V adoption.[16] If the analysis demonstrates superior improvements in NO_x and SO_2 compared to CO, this would support the theory that ineffective policy implementation undermined its expected outcomes. In summary, the central hypotheses are as follows:

1. Compared to a synthetic counterfactual, the adoption of Euro-V standards in Pakistan reduced PM2.5 levels in Karachi.
2. Due to insufficient regulatory enforcement and delays in refinery upgrades, the reduction of PM2.5 levels in Karachi attributable to Euro-V adoption is smaller than anticipated.
3. The slow implementation of the Euro-V policy for domestically produced fuels may be detrimental, as evidenced by comparing NO_x, SO_2, and CO emissions in Pakistan to those of a synthetic Pakistan over time.

3 Methodology: Difference-in-Differences and Synthetic Control Approach

To quantify the effect of the Euro-5 emission standards, a difference-in-differences (DiD) approach was employed. The latter was augmented with a synthetic control to evaluate the impact of a January 2021 Euro-V standard fuel import policy intervention on Karachi's PM2.5 levels. Karachi is treated as the "intervention" city, while a weighted combination of donor cities (which did not implement the policy) serves as a counterfactual. The donor cities – Dhaka (Bangladesh), Dushanbe (Tajikistan), Rangoon (Myanmar), and Baghdad (Iraq) – were selected based on data availability and the absence of similar policy changes during the period of interest. By matching Karachi's preintervention pollution trend to a synthetic composite of these cities, a baseline scenario is constructed representing what Karachi's PM2.5 levels would have been without the policy. This synthetic counterfactual allows the isolation of the policy's effect by comparing Karachi's actual air quality post-2021 to the counterfactual trend.

Monthly PM2.5 concentration data from June 2019 through June 2023 were analyzed, covering approximately 1.5 years before and 2.5 years after the intervention. The intervention date (January 2021) lies roughly in the middle of the time series, ensuring a substantial pre- and post-treatment window for a robust DiD analysis. The DiD was implemented with two-way fixed effects (time and city fixed effects), and an interaction term for Karachi after January 2021 to estimate the treatment effect (Figure 1). A key assumption for validity is that Karachi would have followed the same trend as the synthetic control in the absence of the policy, known as the parallel trends assumption. This assumption was checked by examining pre-2021 trends and conducting a placebo

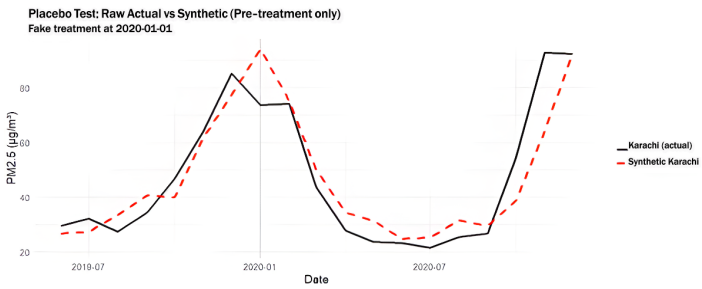

Figure 1: Actual versus synthetic Karachi under a placebo intervention in January 2020. The red, dotted curve represents "Synthetic Karachi", whereas the black, solid curve is "Actual Karachi". The negligible differences in the post-placebo period indicate no effect of the fake treatment.

Below is shown a statistical result of the placebo DiD test between actual and synthetic with fake-treatment. The placebo "effect" is +4.62 g/m³, but it is not statistically significant (p-value = 0.337 > 0.05). This indicates no meaningful divergence between actual and synthetic Karachi at the fake treatment date. The placebo test supports the parallel trends assumption, because the synthetic and actual series do not show a

significant break before the real policy.
No policy effect was detected in any pre-2021 period, confirming that Karachi's PM. tracked the donor-based synthetic closely before the actual intervention. This validation increases confidence that any divergence after 2021 can be attributed to the policy and not a

```
> summary(did_twfe_placebo_st1)
OLS estimation, Dep. Var.: pm25
Observations: 38
Fixed-effects: unit_id: 2, time_id: 19
Standard-errors: IID
Estimate Std. Error t value Pr(>|t|)
treated:post 4.62137 4.67464 0.988606 0.33672
... 2 variables were removed because of collinearity (treated and
post)
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
RMSE: 4.64866 Adj. R2: 0.242642
Within R2: 0.054365
```

3.1 Methodology for Extended Analysis

While the initial analysis focused solely on PM2.5 concentrations to evaluate the impact of Euro-V fuel standards in Karachi, it became evident that examining particulate matter alone was insufficient to capture the full environmental and regulatory implications of the policy. PM2.5 levels reflect only one aspect of air quality, whereas other pollutants, such as sulfur dioxide (SO2), nitrogen oxides (NOx), and carbon monoxide (CO), provide complementary information on fuel combustion efficiency, refinery operations, and compliance with Euro-V standards. To analyze the effects of inefficient implementation of the Euro-V policy, a very similar difference-in-differences (DiD) approach augmented with a synthetic control was utilized. However, due to the availability of data for emissions, country-level data was used. In contrast with the PM2.5 analysis, Pakistan was treated as the "intervention" country, while a weighted combination of donor countries (all of which do not implement the EuroV policy) serves as a counterfactual. The donor countries were Bangladesh, Tajikistan, Myanmar, and Iraq, where the original donor cities were located. This makes it possible to construct a baseline scenario representing what Pakistan's NO_X, SO_2, and CO levels would have been without the Euro-V policy.

4 Data Sources and Measures

The analysis uses compiled monthly PM2.5 data for Karachi and the donor cities, with an emphasis on data quality and comparability across locations. The key outcome measure is the monthly average PM2.5 concentration ($\mu g/m^3$) for each city (Figure 2). Several

4.1 Karachi Data QC

Karachi's reference PM2.5 series was cleaned to ensure it covered the required interval (mid-2019 through mid-2023) with no major gaps and no implausible outliers. The timeframe deliberately includes the January 2021 event and at least two years on either side, capturing seasonal variations and long-term trends (Figure 3). Data points are evenly spaced at one-month intervals; this regular frequency avoids any bias that could arise from irregular sampling. Visual inspection and summary statistics confirmed there were no extreme anomalies in Karachi's data aside from the notable January 2021 spike (which is treated as a real phenomenon rather than a data error).

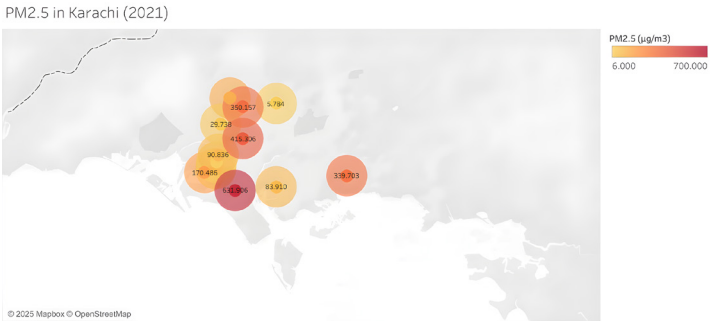

Figure 2: Average PM2.5 levels in Karachi for the year 2021, representing the pre-treatment period. The PM2.5 values range from 6,000 to 700,000. The colour gradient indicates concentration levels, with yellow representing lower values and red indicating higher values. The results for the periods 2019–2020 and 2022–2023 are reported in Appendix (A.1).

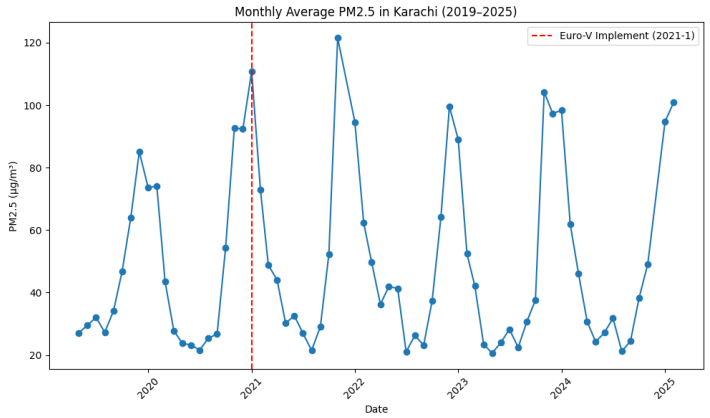

Figure 3: Plot of the PM2.5 recordings for a single station across the period 2019-2025. Note that despite having a visible seasonal pattern for the concentration of PM2.5, no outlier seems to be present. This means that the 2019-2025 Karachi PM 2.5 from a single station recording can be used for our analysis, concluding the QC phase of our research.

4.2 Donor City Data QC

Donor city data were filtered and processed to ensure alignment with Karachi's timeline and conditions. Each donor city's PM2.5 series spans June 2019 to June 2023 without gaps, so that all cities have an overlapping period encompassing the intervention date and a sufficiently long pre- and post-intervention record. Any city that had implemented comparable air quality policies around early 2021 was excluded – this criterion ensures that donor cities provide a genuine control not influenced by the treatment. In the dataset, none of the donors underwent Karachi's January 2021 policy or any analogous measure during the study window. Minor missing data were addressed: the Baghdad series was missing a single monthly value (August 2020), which was imputed by taking the average of the adjacent months (July and September 2020) to maintain continuity. After this correction, all donor series aligned perfectly in time with Karachi's series. This complete and synchronized panel of monthly data enabled the construction of a reliable synthetic control and the execution of DiD regressions without losing observations.

4.3 Donor Selection

The choice of donor cities was guided by both geographic relevance and independence from the policy (Table 1). Focus was placed on large cities in South/Central Asia that suffer from air pollution, but did not implement the specific intervention that Karachi did. The final donor pool (Dhaka, Dushanbe, Rangoon, and Baghdad) provides diversity in regional climate and pollution sources, which helps in constructing a well-rounded counterfactual. It was observed that adding more donors improved the synthetic model fit; for instance, introducing Baghdad's data after filling its small data gap noticeably enhanced the pre-2021 alignment between Karachi and the synthetic PM2.5 trend. This indicates that each donor contributed incremental information to better capture Karachi's baseline pattern.

City	Karachi	Dhaka	Dushanbe	Yangon	Baghdad
Euro-V	Yes	No	No	No	No
Population	18.8 mil	10.2 mil	1.2 mil	5.2 mil	7.9 mil
Port City	Yes	Yes	No	Yes	No
Avg. Temperature	80.8°F	78.8°F	59.4°F	81.3°F	74.3°F
Rain (days/year)	19	118	69	137	34

Table 1: Comparative Characteristics of Karachi (Treated City) and Donor Pool Cities. The table summarises key differences and similarities in policy implementation, population metrics, and climate variables used in the analysis.

4.4 Normalization

PM2.5 concentrations were not normalized by city population. The reasoning is that PM2.5 is an ambient concentration ($\mu g/m^3$) rather than a quantity that scales directly with population. Dividing by population would not yield a meaningful "per capita pollution" metric, since a city's population size does not linearly determine its pollution concentration (many other factors like traffic volume, industry, geography, and weather play roles). Moreover, population differences between Karachi and the donor cities, as well as population growth over the analysis period, were relatively modest. For example, Karachi's metropolitan population was about 16.1 million in 2020, growing to roughly 17.2 million by 2023 (Macrotrends, 2025). Dhaka grew from around 20.3 million in 2019 to 23.2 million in 2023, approximately a 3% annual growth rate (Macrotrends, 2025). Other donors saw similarly small annual population increases on the order of 2–3%. Such changes are not significant enough to warrant adjusting the air quality data, and doing so could actually obscure real pollution trends. In environmental analyzes, per-capita normalisation is more common for gas emissions (e.g., CO_2 or NO_x) rather than for ambient pollutant concentrations. Given these points, the raw PM2.5 concentration data were used for all cities, assuming that any population-related effects on pollution are implicitly captured in the pollution measurements themselves.

4.5 Data Sources for Extended Analysis

To complement the analysis of ambient PM2.5 concentrations, it was necessary to examine emissions data for NOx, SO2, and CO. This analysis also used compiled monthly emissions data for the three target substances. Data from the Emissions Database for Global Atmospheric Research (EDGAR)17, a European Union data source, was compiled and processed. We analyzed the monthly emission data for the three substances from 2008 to 2022, allowing for pre-treatment and post-treatment analysis of the policy intervention (Figure 4). This was the window of time for which data was available for monthly data - this monthly data was necessary to construct a working regression. The data source contained data for each country globally, with every type of emission. To process this data to be used for analysis, the target countries had to be isolated, and all the emission types had to be consolidated for each emission substance.

The Python code in the following page shows the cardinal steps to isolate the target countries only and to sum emissions over all types[18].

```
# ISO3 codes:
# Isolating only countries from synth
target_iso3 = ["PAK", "TJK", "IRQ", "MMR", "BGD"]

# Filter the DataFrame based on country and year
df_filt = df[
    (df["Country_code_A3"].isin(target_iso3)) &
    (df["Year"] >= 2008) &
    (df["Year"] <= 2022)
].copy()

print("\nCountries in filtered data:", df_filt["Country_code_A3"].
    unique())
# Sum over all types
month_cols = ["Jan", "Feb", "Mar", "Apr", "May", "Jun",
              "Jul", "Aug", "Sep", "Oct", "Nov", "Dec"]

# Substance should be NOx (Note: Assuming df_filt has this column)
group_cols = ["Country_code_A3", "Name", "Year", "Substance"]

# Aggregate monthly data by summing across the month columns
df_monthly_totals = (
    df_filt
    .groupby(group_cols, as_index=False)[month_cols]
    .sum()
    .sort_values(["Country_code_A3", "Year"])
)

print("\nMonthly totals (wide) head:\n", df_monthly_totals.head())
```

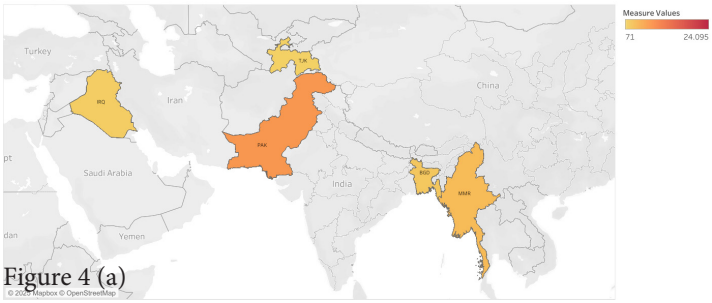

Figure 4 (a)

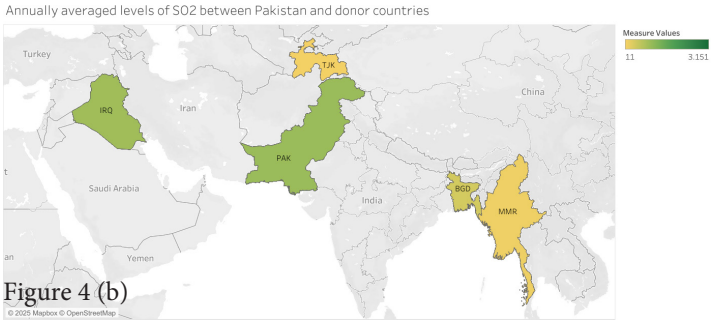

Figure 4 (b)

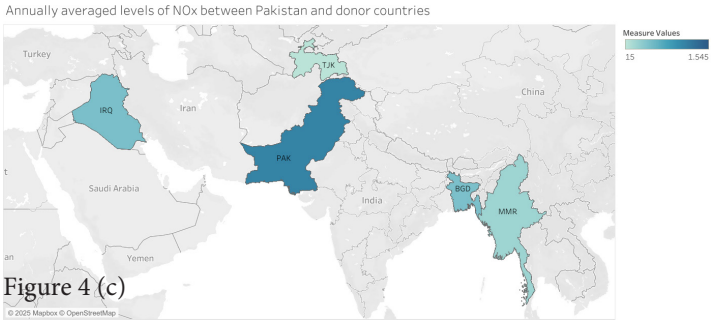

Figure 4 (c)

Figure 4: Annual average levels of CO, SO2, and NOX emissions in Pakistan compared to donor countries from 2008 to 2022.

4.6 Rejection of STL Adjustment in Extended Analysis

Following a similar approach to the seasonal adjustment applied previously, the seasonaltrend decomposition using STL was performed for the three additional substances identified for analysis. This procedure isolates seasonal patterns and trends in the data, allowing for a clearer assessment of the policy's impact on NO_x, SO_2, and CO emissions.

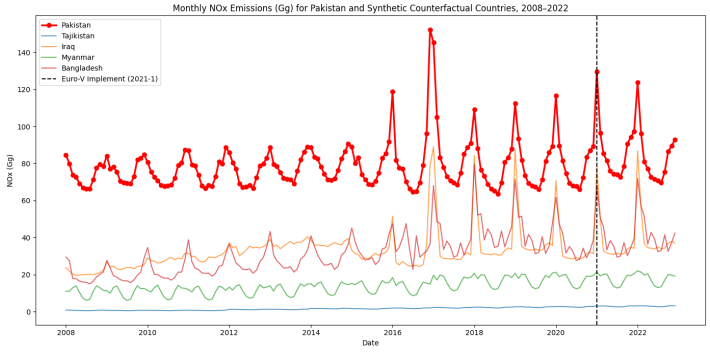

Figure 5: Monthly NOx Emissions for Pakistan and Synthetic Counterfactual Countries, 2008–2022. The dashed vertical line indicates the implementation date of the Euro-V emission standard in 2021.

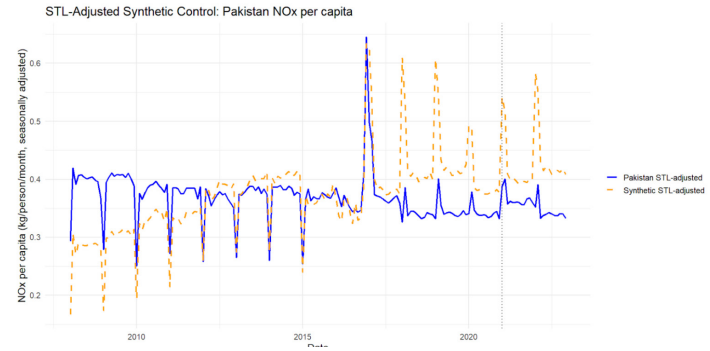

Figure 6: STL-adjusted NOx results illustrating the methodological breakdown caused by data anomalies.

The initial evaluation of the data revealed several noteworthy observations. In particular, a pronounced anomalous increase in NO_x emissions was observed during the 2016–2017 period (Figure 5). Although the general upward trend in emissions since the baseline year of 2008 can largely be attributed to increased industrialisation, this acute, short-term anomaly presents a methodological complication. Employing Seasonal-Trend decomposition using Loess (STL) for seasonal adjustment would likely overcompensate for this substantial spike, producing artificially large negative values in the resulting time series (Figure 6). Accordingly, the subsequent analysis will utilise the non-STL-adjusted data. For this preliminary investigation, this approach offers a more robust and reliable basis for evaluating the impact of the variable of interest.

4.7 Normalization in Extended Analysis

A second observation was the necessity of normalisation (Figure 7). In this case, per-capita normalisation was selected as the preferred approach. Given that Pakistan exhibits substantially higher emissions alongside a correspondingly large population, this step was essential to accurately interpret the data and to construct a synthetic counterfactual that satisfies the pre-implementation trend check. To perform the percapita normalisation, a spreadsheet containing the annual population of each country was merged with the original dataset, after which the emissions data were normalized using information from both sources[19].

```
# Per-Capita Normalisation

# Define the list of columns representing monthly data
month_cols = ["Jan", "Feb", "Mar", "Apr", "May", "Jun",
              "Jul", "Aug", "Sep", "Oct", "Nov", "Dec"]

# Example: NOx in kg per person (1 Gg = 1e6 kg)
for m in month_cols:
    # This loop calculates the per-capita value for each month 'm'
    # df[m] is assumed to be in Gigagrams (Gg).
    # It converts Gg to kg (multiplying by 1e6) and then divides by the
        population.
    df[f"{m}_kg_per_person"] = df[m] * 1e6 / df["population"]
```

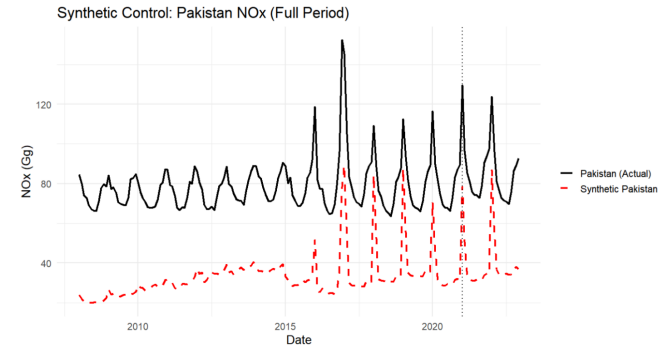

Figure 7: Synthetic Control Results: Actual Pakistan (Solid) versus Synthetic Counterfactual (Dashed) for Total Monthly NOx Emissions. The data are presented before per-capita normalisation.

5 Results and Discussion

5.1 Discussion for PM2.5

The analysis reveals a clear decline in Karachi's PM2.5 levels following the January 2021 intervention, compared to the trajectory the city would likely have had without the policy. The synthetic control closely reproduced Karachi's actual PM2.5 trend in the pre-2021 period, which bolsters confidence in its validity as a counterfactual. Starting immediately after January 2021, Karachi's observed PM2.5 drops below the synthetic
counterfactual curve, indicating an improvement in air quality relative to the no-policy scenario. In quantitative terms, the DiD estimation with two-way fixed effects finds an average PM2.5 reduction of about 11.7 µg/m3 attributable to the policy in the post-intervention period. This drop is statistically significant ($p \approx 0.008$), supporting the hypothesis that the intervention had a material beneficial effect on air quality. In practical terms, a double-digit decrease in µg/m3 is substantial. Karachi's annual average PM2.5 was around 40 µg/m3 in recent years, so a ~11.7 µg/m3 improvement represents a considerable air quality gain. This result confirms the first hypothesis that the policy implemented in Karachi has led to a measurable and significant reduction in fine particulate pollution.

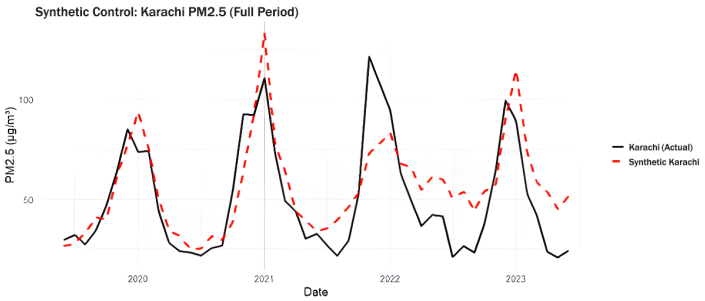

Figure 8: Comparison of actual PM2.5 levels in Karachi (black) with the revised synthetic Karachi (red), constructed using four donor cities.

A statistical result of the DiD test between actual and synthetic PM2.5 in the posttreatment period was obtained using OLS with a Two-Way Fixed Effects model estimation (Figure 8). The estimated treatment effect (Euro-V on Karachi PM2.5) is –11.67 µg/m3, indicating that, after the intervention, actual Karachi PM2.5 is 11.7 units lower relative to its synthetic control. This effect is statistically significant at the 1% level, with p-value = 0.008.

```
> summary(did_twfe)
OLS estimation, Dep. Var.: pm25
Observations: 96
Fixed-effects: unit\_id: 2, time\_id: 48
Standard-errors: IID

            Estimate  Std. Error  t value  Pr(>|t|)
treated:post  -11.6652   4.22766  -2.75925  0.0082882  **

... 2 variables were removed because of collinearity (treated and
    post)
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
RMSE: 7.01106     Adj. R2: 0.841977
Within R2: 0.142006
```

While investigating the PM2.5 trend in Karachi, a pronounced down peak in December 2020 and an up peak in January 2021 emerged when STL (Seasonal-Trend decomposition using 'Loess' method) was applied in R, which was intended to isolate the target policy effect by removing any seasonal repeating patterns (Figure 9).

```
STL: Seasonal Adjustment

ts_actual    <- ts(results$actual,    frequency = 12)
ts_synthetic <- ts(results$synthetic, frequency = 12)

stl_actual    <- stl(ts_actual, s.window="periodic", robust =
   TRUE)
stl_synthetic <- stl(ts_synthetic, s.window="periodic", robust =
   TRUE)

# Seasonally adjusted (remove seasonal component)
results$actual_stl    <- as.numeric(ts_actual    - stl_actual$
   time.series[,"seasonal"])
results$synthetic_stl <- as.numeric(ts_synthetic - stl_synthetic$
   time.series[,"seasonal"])
```

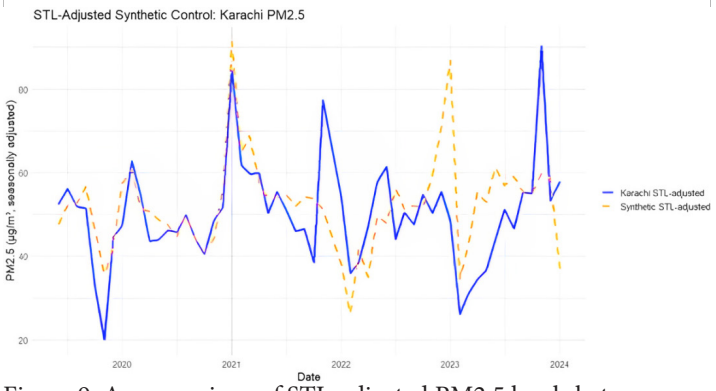

Figure 9: A comparison of STL-adjusted PM2.5 levels between actual Karachi and synthetic Karachi.

These abnormal spikes were not visible in the raw data comparison between actual Karachi and the synthetic control, which instead showed a typical seasonal fluctuation consistent across years. This inconsistency indicates that the STL decomposition introduced an artifact, possibly due to several reasons such as the short data window (only 20 months for pre-treatment), which limits STL's ability to estimate stable seasonal components, the abrupt policy effect around the treatment point (Euro-V fuel adoption), which STL may misinterpret as irregular seasonality, and asymmetric seasonal patterns in PM2.5 levels across cities, which could violate STL's assumption of regular periodicity. Therefore, it can be concluded that STL should not be applied to this dataset, as it distorts rather than clarifies the true intervention effect. The raw PM2.5 time series comparison between Karachi and its synthetic control already reveals a clear treatment effect without any need for additional decomposition.

While the results strongly indicate a policy effect, some important caveats remain. First, the donor pool is relatively small (four cities). A larger set of control cities could potentially yield an even more accurate synthetic counterfactual and bolster the robustness of the findings. With limited donors, there is a risk that idiosyncratic factors in any one donor city (e.g., an unobserved local policy or an unusual pollution event) could influence the synthetic composite. This was mitigated by carefully choosing donors unlikely to have their own interventions and by testing the model's sensitivity (for example, adding the fourth donor, Baghdad, improved the pre-policy fit noticeably, which suggests the donor set was moving in the right direction). Still, the generalizability of the exact numeric estimate (-11.7 µg/m³) could be affected by the small donor sample; incorporating more cities or data (if available) would strengthen confidence in the exact magnitude of the effect.

Second, although the parallel trend assumption appears to hold in the pre-intervention period, not all confounding factors can be ruled out. Differences in climate and geography between Karachi and the donors mean that certain pollution drivers (e.g., monsoonal patterns or dust events) are not perfectly replicated by the synthetic control. It is assumed that these influences average out via the donor combination, but some residual differences could remain. Lastly, the analysis focuses on the immediate PM2.5 metric and does not delve into other pollutants or health outcomes – the policy likely had broader impacts (positive or negative) not captured here Despite these limitations, the evidence strongly suggests that the January 2021 clean-air policy intervention in Karachi has been effective. The significant drop in PM2.5 and the sustained lower levels post-2021, relative to the counterfactual, underscore the value of such environmental policies in combating urban air pollution and improving public health in mega-cities like Karachi.

5.2 Discussion for NO_x, SO_2, and CO

The results in Figure 10 indicate that both NO_x and SO_2 exhibit a clear post-treatment effect. Although this effect is not perfectly visible in the graphical representations—largely due to imperfect pre-treatment matching—it becomes evident through the DiD summary statistics. CO, however, behaves differently: its estimated effect is positive, meaning Pakistan's emissions increased relative to its synthetic counterpart. This outcome directly contradicts the third hypothesis, as CO did not display a more substantial emissions reduction than NO_x or SO_2. In fact, the Carbon Monoxide results were not statistically significant at all (Table 2).

This may suggest that the relative difference between imported petrol and domestically refined diesel does

not meaningfully influence CO emissions. A more plausible explanation, however, is that changes of this nature are simply not statistically detectable for CO. In Pakistan, CO originates from numerous additional sources—most notably widespread biomass burning. Moreover, the adoption of Euro-V fuel standards does not imply that the entire vehicle fleet complies with corresponding technological or maintenance standards. Many older or poorly maintained vehicles continue to emit substantial quantities of CO regardless of fuel quality. Compounding this issue, not all fuel in Pakistan enters the market through legal channels and is therefore not subject to regulatory oversight. The fuel black market remains extensive, with an estimated 5–6 million litres smuggled across the Iranian border each day.

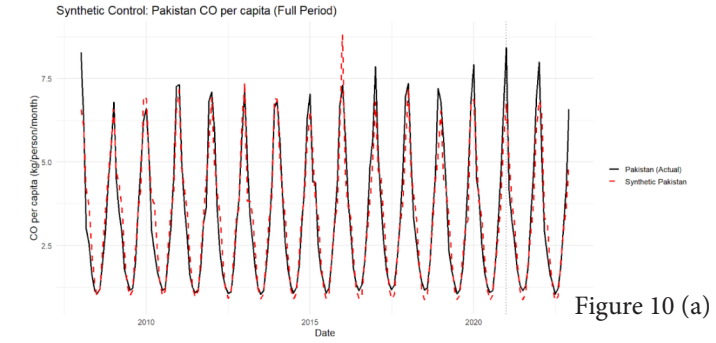

Figure 10 (a)

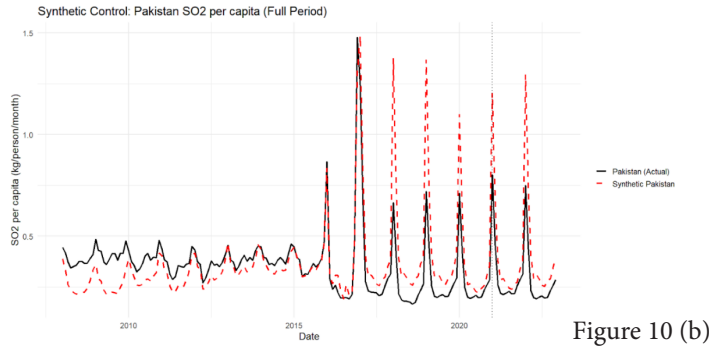

Figure 10 (b)

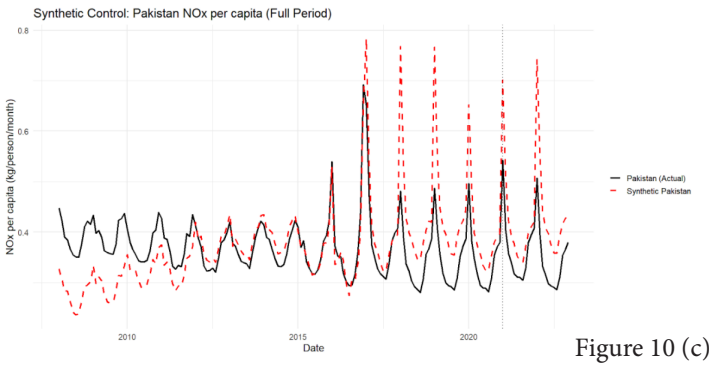

Figure 10 (c)

Figure 10: Synthetic Control Results: Monthly CO, SO_2, NO_X Emissions Normalized by Population (kg/person/month). The plot compares Pakistan's actual emissions (Solid) against the Synthetic Counterfactual (Dashed), demonstrating a strong pre-intervention common trend.

Overall, it remains difficult to quantify the precise negative impact of Pakistan's delayed refinery upgrades on national air quality and emissions. Nevertheless, given the policy's demonstrated positive effect on PM2.5 concentrations in Karachi and the statistically significant improvements observed for NO_x and SO_2, it is evident that the lack of modernisation in domestic refineries continues to hinder Pakistan's environmental progress.

Pollutant	DiD	p-val	Significance
NO_x	-0.122 kg/p/m	5.89×10^{-7}	Yes
SO_2	-0.0758 kg/p/m	1.90×10^{-5}	Yes
CO	$+0.1941$ kg/p/m	0.18	No

Table 2: DiD estimates for the three pollutants.

5.3 Policy Implications

The analysis of PM2.5 trends in Karachi, in combination with the evaluation of secondary emissions data (CO2, NOx, SO2) and the synthetic control framework, provides several key policy implications for Pakistan's air quality management.

1. The drastic decline in the levels of PM2.5 after the introduction of the Euro-V fuel standards in January 2021 is evidence that regulations can achieve measurable environmental dividends, which makes further upgrading towards Euro-VI standards a logical next step. It would mean another stride toward a more stringently controlled fuel quality and, perhaps, heightening the declines of fine particulate pollution across the urban centres.
2. The incomplete adoption and delayed refinery upgrades have limited the full effectiveness of the policy to date, and point to a need for stronger regulatory enforcement. The government may consider stricter penalties, compulsory deadlines for compliance, or even shutdowns for non-compliant refineries to ensure that all fuel supplied within the country meets the standards. Comparing imported fuels, already at the Euro-V standards, with domestically refined diesel, which is producing disproportionately higher NO_x and SO_2 emissions, suggests that incomplete implementation in the refineries at home undermines the expected gains from such fuel policies. Policymakers should, thus, make refinery modernisation a priority and assess interventions on emission hotspots associated with outdated infrastructure.

3. Scientific evidence testifies that outdated emission standards need urgent attention to prevent wider public health impacts and boost climate resilience.

All together, these findings indicate that a strategy coordinating enhanced enforcement with infrastructural improvements and progressive enhancements of fuel standards will be imperative for the sustainability of long-term air quality improvement in Karachi, indeed across all urban areas of Pakistan.

5.4 Policy Recommendations

1. **Transitioning directly to Euro-6 fuel standards**:[20] Pakistan's government should expand the ongoing refinery modernisation for Euro-5 and implement a complete transition to Euro-VI without accepting an inferior interim standard. The completion of both upgrades in one infrastructure cycle will avoid capital duplication, decrease downtime, and ensure compliance with global emissions standards over the long term. In this regard, the revised national deadline, extended by about two years, would give refineries a realistic timeframe for ordering machinery, undertaking technical revamping, and arranging finances.
2. **Strict compliance with Euro-5 by 2027 if the transition to Euro-6 is rejected**: This should involve the establishment of quarterly audits for compliance, clear public reporting, and legally binding penalties in the case of non-compliance, including suspension of production or monetary fines. From 2027 onwards, all sub-Euro-5 fuel must be banned from legal sale, and distribution networks must be compelled to verify fuel quality at each point in the supply chain.
3. **Strengthening and modernisation of national vehicle emissions standards**: For cleaner fuel to translate into actual emission reductions, Pakistan needs to put in place strict controls for new and existing vehicles. This would cover the introduction of Euro-5 or Euro-6 emission standards for all newly imported or locally manufactured vehicles; carrying out mandatory annual vehicle inspections for emission compliance; and implementing a phase-in retirement program for high-emitting vehicles, especially the older diesel buses, rickshaws, and freight trucks.
4. **Deterrence of fuel smuggling and the black-market distribution system**: Pakistan needs to extend border-security operations along the Balochistan-Iran corridor, where approximately 5-6 million litres are smuggled every day. This would include deploying additional surveillance technology, targeted enforcement operations, and coordination in action by customs, provincial governments, and the Frontier Corps. Adding traceability markers to legally refined fuel and increasing penalties for illegal resale would also considerably decrease the circulation of unregulated, high-sulfur fuel.`

Acknowledgements

We sincerely express our gratitude to Prof. Dr Asensio and Ms Lengefeld for their thoughtful guidance and feedback throughout the writing process in the context of their PUBL 3042 class. We thank Ms Davis for her valuable support on Tableau Cloud. We extend our appreciation to the PAQI (Pakistan Air Quality Initiative) for the provision of the datasets.

Author Contributions

Conceptualization: D.G., J.K., M.P.; Data curation: J.K.; Formal Analysis: D.G., J.K; Investigation: D.G., J.K., M.P.; Methodology: D.G., J.K., M.P.; Resources: D.G., J.K., M.P.; Software: D.G., J.K.; Validation: M.P.; Visualization: M.P.; Writing – original draft: D.G., J.K., M.P.; Writing – review and editing: M.P.

2019–2020, 2022-2023

A Appendix
A.1 Average PM2.5 levels in Karachi for the year

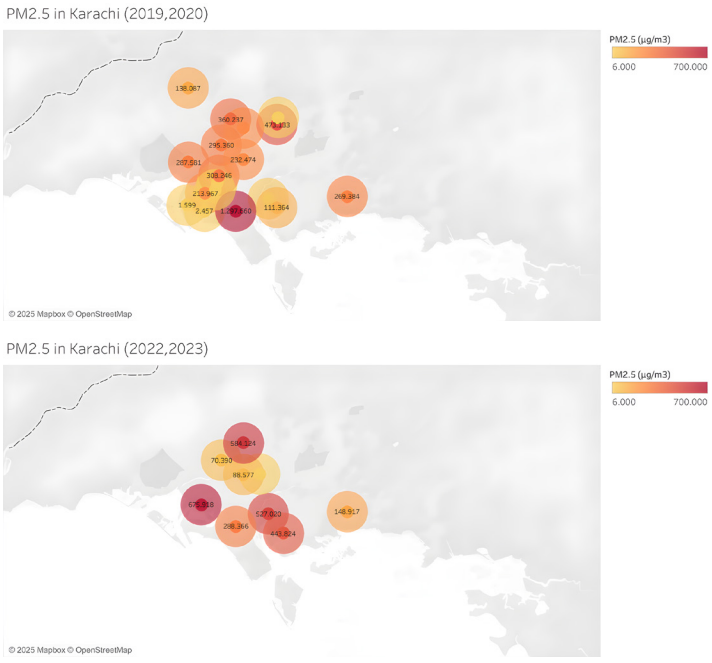

Figure 11: Average PM2.5 levels in Karachi for the year 2019-2020 and 2022-2023, representing the pre-treatment period. The PM2.5 values range from 6,000 to 700,000. The color gradient indicates concentration levels, with yellow representing lower values and red indicating higher values. The results for the period 2021 is shown in Figure 3.

Access references for this paper using the QR code below.

Editor Biographies



Melody Lee

Melody is a second-year undergraduate student studying Computer Science, with minors in math and physics. Her research is mainly in quantum algorithms for applied science.



Kathy Peng

Kathy is a second-year neuroscience major and CS minor. She is very passionate about the intersection of computer science in neuroscience research and its potential in improving ethical and sustainable practices!



Ananya Pasunuri

Ananya is a sophomore, mechanical engineering major with an aerospace minor. She has a passion for research within the aerospace and sustainability industry!



Hanna Zhang

Hanna is a first-year, industrial engineering major with a minor in robotics. Shes interested in the impacts of natural disasters on communities and heathcare systems.



Zoe Chow

Zoe is a first-year, biomedical engineering student. She is interested in both research and design, and loves staying up to date with the advancements Georgia Tech students have been able to achieve.



Tvishi Ahluwalia

Tvishi is a second-year biomedical engineering student with interests in immunology and drug delivery. She is currently conducting research on targeting immunosuppressive cells to improve healing.



Shreya Iyer

Shreya Iyer is a second-year electrical engineering student with a minor in artificial intelligence. She is currently researching the coexistence of non-terrestrial and terrestrial networks to advance satellite-enabled positioning, navigation, and timing (PNT) systems worldwide.

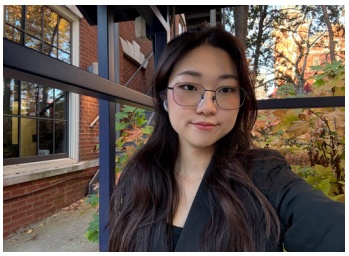

Yeaeun (Caroline) Song

Caroline Song is a third-year Biomedical Engineering student at Georgia Tech. Her research focuses on medical imaging AI, including deep learning applications for Parkinson's disease diagnosis. Originally from South Korea, she spends her free time playing cello and crocheting.

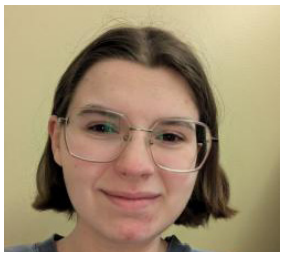

Olivia Lohr

Olivia Lohr is a third-year Literature, Media, and Communication major with minors in biology and science communication and policy. She holds a background in technical communication and public outreach, with past research work studying perceptions of leadership in academic settings.



Diya Sharma

Diya is a Computer Science major with a minor in Neuroscience, specializing in AI and Human-Computer Interaction (HCI). She conducts bio-translational wearable device research at the BITN Lab.



Mirielle Dogini

Mirielle is a Biology student with a minor in Health and Medical Services. She is interested in regenerative medicine and tissue engineering research, and she's currently in a lab that is working on wound repair in children. Mirielle likes to run, dance, and paint in my free time.



Isabel Hollhumer

Isabel is a biomedical engineering major. She is currently conducting research on cystic fibrosis treatments and innate immune responses to tuberculosis. Additionally, she leads three international research projects in sub-Saharan Africa.

ART · PLUMBING
SERVICE
·EST 1919·

Over "100" years in the Plumbing Service Industry





- 24/7 Plumbing Service & Maintenance "365 Days"
- 24/7 emergency response to any system or utility failure
- General repairs – regardless of size!
- TV/Video pipe inspection
- Back flow prevention upgrades, repairs and state certification
- Utility pipe leaks (water, gas and sewer)
- Water jet sewer cleaning (up to 30" in diameter)
- Sewer, drain and grease trap cleaning
- Fire hydrant repair
- Process piping
- Water heaters
- Sewer ejectors-repair & preventative maintenance
- Natural gas
- LEED Buildings; Water Conservation and Water Metering

"We are here to solve your challenges"



Phone: 678-486-2540 • Fax: 770-433-1452
1847 South Cobb Industrial Blvd. Smyrna GA 30082 · www.artplumbing.com

LADE
ELECTRIC SUPPLY
SINCE 1973

Midtown Branch
204 14th St. NW
Atlanta, GA 30318
Office: 770.278.9222
WWW.LADESUPPLY.COM

M.E. Contractors, Inc.



1263 Ellsworth Drive NW, Atlanta, GA 30318
Phone: (404)-351-2803
www.mecontractorsinc.com

FACILITY CONTROLS GROUP

Best Wishes to the Students of GA Tech!

Jeff Willis
President

5174 Hathburn Court
Dunwoody, GA 30338
Ph: 770-633-3242
jeffwillis102@hotmail.com

What is *The Tower*?

The Tower is the only publication of the Georgia Institute of Technology that publishes peer-reviewed undergraduate research from all academic disciplines. A biannual publication, The Tower provides a comprehensive forum for discourse on cutting-edge research. In order to ensure only high quality research is published, all submissions undergo stringent blind, peer-review.

Papers are reviewed based upon methodology, originality, and relevance of results. Statistical analysis should be used and reported where appropriate. Depth of analysis should be comprehensive. The most important consideration is to ensure that the paper is written for a general, educated audience. If the final set of accepted articles is not balanced to the satisfaction of the Editorial Board, accepted articles may be "held over" until the next year's issue. Additionally, The Tower's editorial staff maintains the right to revise articles for clarity and modifications for layout adherence.

Who Is Eligible?

To submit, students must:

1. Be undergraduate students currently enrolled at the Georgia Institute of Technology.

2. Submit research work done while enrolled as a student at the institute.

3. Not be graduating in the fall semester of the year of submission.

How Do Submissions Work?

Submissions are reviewed by trained graduate and undergraduate students as well as expert faculty in the field of the submission. The peer-review process takes several few months, on average. Often manuscripts will need several cycles of review and revision before they can be accepted for publication. Therefore, it is highly suggested that every effort is made to be early and prompt with submissions to ensure inclusion in the next issue of the journal.

1. Submit your manuscript

Submissions open in late fall until early January—check out the website for more updates.

2. Notification of selection

Authors of selected papers will be notified in the weeks following the submission deadline.

3. Collaboration with editors

Authors will be connected with an editor from *The Tower* and a graduate student or faculty editing advisor to revise and finalize the manuscript. Authors are expected to keep their editors updated every two weeks.

4. Final Publication

The manuscript will be formatted into the final journal. Journals will be printed and distributed by late spring.

THE TOWER