NYU | TANDON SCHOOL OF ENGINEERING









2026

UNDERGRADUATE SUMMER RESEARCH PROGRAM

STUDENT ABSTRACT BOOKLET

A MESSAGE FROM THE VICE DEAN

The NYU Tandon School of Engineering's Undergraduate Summer Research Program (UGSRP) provides a hands-on research opportunity for students across NYU Tandon, NYU Abu Dhabi, NYU Shanghai, NYU College of Arts and Science, NYU-KAIST, and other select institutions. During this 10-week program, student researchers collaborate with faculty mentors, graduate scholars, and research staff on high-impact projects. By connecting classroom learning with real-world challenges, UGSRP expands academic horizons and prepares participants for lasting scholarly success.

Guided by faculty across Tandon's Impact Areas, participants work to build smart, secure, sustainable, and healthy urban communities. Faculty and research supervisors provide essential guidance in research methodologies and field-specific practices. Alongside lab work, students attend lectures and professional development workshops. Special thanks to Matthew Frenkel for leading the Student2Scholar (S2S) Series, and to our guest speakers for sharing valuable guidance for our students' futures.

Financial support from NYU Tandon, faculty research grants, and external donors is vital to the program. This year marks the 15th anniversary of the Thompson Bartlett (TB) Fellowship, which supported ten female researchers thanks to Mrs. Dede Bartlett. Her father, George Juul Thompson, graduated from the Electrical Engineering program at the Polytechnic Institute of Brooklyn in 1930; we sincerely thank Dede for her ongoing support. Other fellowships include the Tandon Scholars Fellowship, the Visiting Student Fellowship, and the NYU-American Heart Association (AHA) Cardiovascular Engineering Undergraduate Summer Research Fellowship, featuring projects with NYU Langone faculty. Donors' generous contributions allow us to expand opportunities for students and mentors alike.



This summer offered both on-campus and remote research opportunities, creating a hybrid experience that mirrors today's professional landscape. The program concluded with a research poster session celebrating student achievements. The abstracts and posters in this booklet highlight the diverse research conducted this summer.

Coordinating this program requires significant effort. Special thanks to Sara-Lee Ramsawak, Assistant Dean of Undergraduate Academics and Global Engagement, who oversees the program and manages booklet production. Miri Cha helped compile this year's Abstract Booklet. Jen Piro, Director of Undergraduate Programs, managed student placements and stipends, supported by Raffy Orbe on application enhancements. Kat Arrendondo, Manager of Undergraduate Programs, led the event calendar and engagement initiatives. Thank you to the entire team!

Congratulations to all participating student researchers! Thank you to our mentors, staff, Dede Bartlett, Matt Frenkel, coordinators, and guest speakers. I am extremely proud of this program and look forward to future summers of scholarly growth.

Peter J. Voltz

Peter Voltz, Vice Dean for Academic Affairs

A MESSAGE FROM THE COORDINATORS

The Tandon Undergraduate Summer Research Program (UGSRP) was initiated at NYU Tandon in 2007. The Office of Undergraduate Academics (UGA) has coordinated the program since 2011, driving its growth from 60 student participants to over 130 this year.

Our program encompasses a diverse group of faculty and research projects across NYU Tandon academic departments, NYU Wireless, the NYU Center for Urban Science and Progress, NYU Langone, and various VIP projects.

We have worked tirelessly to match students with best-suited project supervisors and research staff members to foster fruitful research and strong mentorship. In addition to hands-on research, we offer dedicated academic, career, and personal development seminars, lectures, workshops, and events designed to deepen students' research knowledge and highlight future professional opportunities and paths.

Our commitment to developing, maintaining, and continually enhancing this program remains steadfast. We extend our sincere gratitude to everyone who joined and supported this summer's program to ensure its seamless execution. We congratulate all student participants on a successful summer and hope you enjoyed the experience!



Sara-lee Ramsawak
Assistant Dean of Undergraduate Academics & Global Engagement



Jen Piro
Director of Undergraduate Programs



Kat Arredondo
Manager of Undergraduate Programs



Miri Cha
UGA Graduate Assistant
UGSRP Booklet Designer



Raffy Orbe
UGA Graduate Assistant
UGSRP Application Developer

Commemorating 15 Years of the Thompson Bartlett Fellowship Program:
Empowering Undergraduate Research at NYU Tandon



Since 2007, each summer members of the Tandon School of Engineering faculty with yearly expansion to include several other offices, centers, and labs have opened spots on their research projects to allow rising sophomore, junior, and senior level undergraduate students to work at a level generally reserved for master's and doctoral candidates through the Undergraduate Summer Research Program. In addition to the work they do on the projects, students participate in programming coordinated by the Undergraduate Academics Office.

Since 2012, numerous outstanding female undergraduates have been eligible for an additional honor: the Thompson Bartlett Fellowship, sponsored by longtime benefactor Dede Bartlett. Bartlett — a former Fortune 25 executive with Exxon Mobil and Altria, philanthropist, advocate for victims of domestic abuse, and 2025 NYU Tandon commencement speaker — launched the fellowship in honor of her father, George Juul Thompson. Thompson graduated in 1930 from the Polytechnic Institute of Brooklyn, received naval awards for his service, and served as an electrical engineering instructor at the school from 1948 to 1968. "During his two decades of teaching, he had only one female student," Bartlett explained, "and I wanted to help change that." Bartlett, who has also supported internship programs at Vassar and NYU, established the fellowship to create tangible research opportunities for women in engineering.

Summer 2026 marks the 15th anniversary of the fellowship program. Initially sponsoring 3 students, Dede Bartlett and her son, Jim Bartlett, have expanded the program to annually support 10 students. The fellowship forms part of a structured 10-week paid summer research assistantship program open to undergraduate students across all majors, enabling hands-on research and mentorship.

The program boasts significant Student and Alumni Impact: to date, approximately 138 students have been sponsored, with about 70% pursuing careers in STEM fields and roughly 45% earning or pursuing advanced degrees. Alumni have secured prominent career placements and admissions to top graduate institutions. At the 15-year celebration event, alumni panelists such as Sarah Appelbaum and Rana Mohamed shared their professional journeys, reflecting on how the fellowship helped shape their careers in STEM and supported their technical and leadership growth.

The Undergraduate Academics Office, the Undergraduate Summer Research Program, and the New York University Tandon School of Engineering are extremely thankful to Dede and Jim Bartlett for their contributions and support of our female students. Over the years, a large network of exceptional women have developed and have been nurtured through this Fellowship. We are, and continue to be, impressed with them and all that they do. We look forward to the future achievements of all our Thompson Bartlett alumni!

THOMPSON BARTLETT FELLOWSHIP CELEBRATES 15 YEARS OF IMPRESSIVE WOMEN













Contents

Project Supervisors & Research Staff

Calendar of Events

UGSRP Fellowships

Areas of Excellence

Industrial, Urban & Environmental Sustainability

Aaryan Vira
Anna Schwarz
Cleatus Oakes
Daksh Mehta
Daniel Zheng
Darren Sea
Gaurav Poddar
Jed Chen
Jed Liu
Hyunggeun Yoon
Ian Yum
Jonathan Platt
Kelsey Cheung
Luyu Zhao
Maxim Belykh
Morgan Huzjak
Muhammet Mustafa Diri
Nicholas Reinoso
Nour Al Saqqa
Reet Barua
Shely Dash
Sienna Suarez
Xiyou Liao
Yixia Yu
Yunyi Wang
Zhijian Mao

Quantum Science & Technologies

Cayden Auyeung
Cooper Stephenson
Kayla Vu
Koki Furusho
Ryan Kim
Sophia Wang
Sungyun Sohn

Robotics & Embodied Intelligence

Alexandre Aarao Reis Arturi
Allan Zhang
Aman Sunesh
Demarce Williams
He Tu
Huiyi Yang
Ian Fiterman
Junho So
Liuyang Bai
Michael Carcamo
Nazmus Sakib
Noah Mays-Smith
Patrick Hsu
Ruoshui Li
Sabrina Addo
Tejas Siddaramaiah
Xiangxiang Hu
Xingtong Liu
Xufei Zhang

Engineering Health

Abigail Huang
Aiden Li
Anika Das
Anjana Kottaiveedu
Arushi Desai
Barbara Czarkowska
Cansu Altintas
Cayden Auyeung
Cecelia Meunier
Chiara Pinci
Cindy Chen
Cristiana Pintea
David Chauca
David Kedem
Evalia Koumpourlis
Fanni Miettinen
Fei Wu
I-Ting Chen
Jessie Huang
Jing Hu
Kumael Mohib
Landon Nguyen
Loubna Kahramane
Luhan Zhao
Marissa Gani
Mayuri Patil
Muhammad Zain Naveed
Natalie Nylaende
Ottilie Zachhuber
Oumaima Bouanani Idrissi
Reeda Rashid
Ryan Abdullai
Ryan Zheng
Shangrui Wang
Steffani Rajapaksha
Sulaiman Rasuli
Veer Rushin Choksi
Vicente Correia
Yashvi Rastogi
Yettie Jung
Yichen Zhao
Yoonjae Andrew Joung
Zhaorong Zeng

Systems Engineering & Complex Decision-Making

Anemka Lydia Legbara
Annie Mei
Darryl Linardi
Evan Dworkin
Jonathan Shan
Kaavya Trivedi
Olivia Vikman
Tiankai Zhang
Umar Ali Abbasi
Xinyi Li
Xiyang Qiu
Yucai Li
Ziqi Gu

Engineering & Culture

Eric Tsang Xu
Genesis Tamara Reyes Casas
Irmak Caglar
Kacper Jucha
Madeleine Chan
Mark Ionis

Pinkerton Lab

William Zhang
Evalia Koumpourlis

UGSRP Supervisors

PROJECT SUPERVISORS

Andrea Silverman
Augustin Guibaud
Austin Rovinski
Ayaskanta Sahu
Bertha Jimenez Leon
Bruce Garetz
Chen Feng
Danny Yuxing Huang
David Truong
Dharik S. Mallapragada
Dzung Luong
Eray Aydil
Fraida Fund
Graham Dove
Ingrid Paredes
Irene de Lazaro
Jason Röhr
Jeffrey Hubbell
Jin Montclare
Jin Ryoun Kim
Joo H. Kim
Joseph Chow
Juanita Hidalgo
Katsuo Kurabayashi
Manny Patole
Mark Yarmarkovich
Maurizio Porfiri
Miguel Modestino
Nathalie Pinkerton
Nikhil Gupta
Paola Bardetti
Pavel Kots
Phillip Rauscher
Rakesh Behera
Raz Saremi
Richard A Stein
Rose Faghih
Rui Li
Ryan Hartman
Semiha Ergan
Shivendra Panwar
Tega Brain
Thorsten Kirsch
Torsten Suel
Vedant Das Swain
Weiqiang Chen
Wenqi Cui
Xin Wang
Yao Wang
Yong Liu

RESEARCH STAFFS

Aaron Alpar
Adrienne Fueredi
Ahmadreza Ahmadjou
Alma Rocha
Anand Lalasaheb Jagadale
Caleb Beckwith
Camilla Ancona
Canying Wang
César Arana
Charlie Mydlarz
Dania Muhieddine Orfali
Grace Liu
Heng Quan
Isabella Huang
Ivan Pinos
Jiangce Chen
Jing Zhang
Juexiao Zhang
Laura Pineiro Alonso
Michael Napoli
Minkyung Kim
Ofek Lauber
Paloma Otsuka Kotani
Pooja Shah
Pulkita Jain
Rachel Pollard
Ran Gong
Sameer Upadhye
Shean Fu Phen
Shuo Zhang
Sihang Li
Sumanth Dara
Susanna Jaramillo
Ulises Gutierrez Nunez
Vedant Nadoda
Xinhao Liu
Yongqing Liang
Yujing Song

Calendar of Events

JUNE

1 Program Begins

3 S2S: An Introduction to Engineering Research

4 Translating Research for a Public Audience

9 The Lab Link-Up: Summer Research Social

10 S2S: Finding what you Need: Building a Knowledge Base for Research Questions

11 Navigating Plagiarism in the Age of AI

17 S2S: Communication: How to Effectively Share your Discovers in a Digital World

18 Mentorship Beyond the Lab

24 S2S: Introduction to Visual Communication: Building Effective Posters and Presentations

25 NSF Graduate Research Fellowship Program (GRFP)

JULY

1 S2S: Data Visualization with Python

2 Applying to Research Fellowships / Scholarships

8 S2S: Intro to Research Data Management

9 Graduate School Application Information Session

11 Luna Park at Coney Island with UGA

15 S2S: Intro to Git and GItHub

16 Demystifying the PhD: General Information Session

22 Where will my Research go? Sharing your work through NYU Ultra Violet

23 Reflecting Towards a Purpose

27 Panel Event with Lunch: 15 Year anniversary of Thompson-Bartlett Fellowship

AUGUST

7 UGSRP Symposium & Research Poster Showcase

UGSRP Fellowships

American Heart Association (AHA)
The NYU-AHA Cardiovascular Engineering Fellowship Program is supported by the American Heart Association's institutional award for undergraduate student training. This fellowship provides support for five undergraduate students, offering them the opportunity to engage in cutting-edge cardiovascular engineering research.

Tandon Scholars Fellowship
The Tandon Scholars Fellowship supports high-achieving undergraduates with the potential to become global leaders in innovation, technology, research, and entrepreneurship through the GLASS Scholars Program. This fellowship provides resources, mentorship, and opportunities for upper level GLASS students to engage in cutting-edge research and collaborate with distinguished faculty and industry professionals. It emphasizes developing socially conscious engineering leaders ready to impact society significantly.

Thompson-Bartlett Fellowship
Established in 2012 and generously supported by Dede Bartlett, the Thompson Bartlett Fellowship funds summer research internships for outstanding women undergraduate students across various majors at NYU Tandon. This includes those majoring in engineering fields, as well as those pursuing STEM-related projects.

Visiting Student Fellowship
The Visiting Student Fellowship provides a unique opportunity for select non-NYU students to participate in research at NYU each year. This fellowship aims to diversify the research community by incorporating students from various institutions and backgrounds. Participants gain access to NYU's cutting-edge research facilities, collaborate with leading scholars, and engage in innovative projects across multiple disciplines.

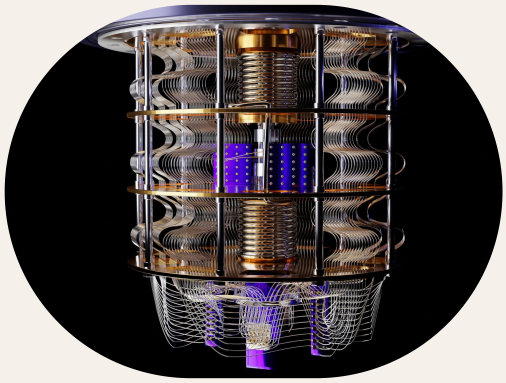
Impact Areas

The NYU Tandon School of Engineering is home to unconventional engineers — a vibrant community of renowned researchers, driven students, strategic industry partners, and ambitious entrepreneurs uniting to solve modern humanity's greatest challenges.

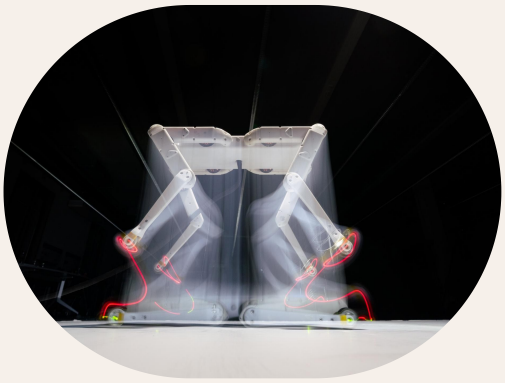
NYU Tandon's research and education are strategically directed toward six cross-disciplinary impact areas:



Industrial, Urban & Environmental Sustainability



Quantum Science & Technologies



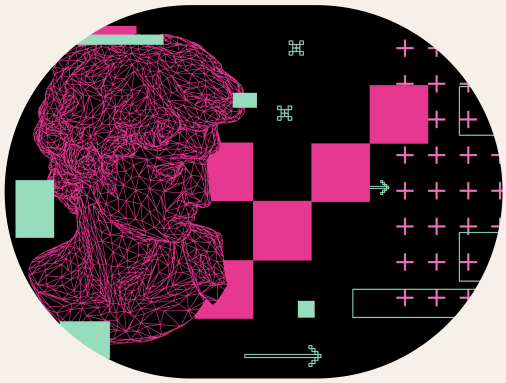
Robotics & Embodied Intelligence



Engineering Health



Systems Engineering & Complex Decision-Making



Engineering & Culture



Industrial, Urban & Environmental Sustainability

NYU Tandon's Sustainable Engineering Initiative develops practical engineering solutions to address climate change, environmental contamination, and societal resilience by advancing technologies in sustainable manufacturing, clean energy integration, smart buildings, water sustainability, AI-enhanced environmental modeling, and materials circularity. Leveraging innovations in machine learning and automation, the initiative accelerates the development of environmentally sustainable technologies while maintaining economic value and performance. Through interdisciplinary collaboration across engineering, science, law, business, and medicine, as well as partnerships with New York City's innovation ecosystem, Tandon integrates technological advancement with risk management, entrepreneurship, and policy to create solutions that protect environmental and human health while promoting long-term economic prosperity.

Industrial, Urban & Environmental Sustainability



Aaryan Vira
Rice University
Spring 2029
Mechanical Engineering, B.S.

Techno-economic Assessment and Sensitivity Analysis of Ammonia as a Liquid Hydrogen Carrier for Long Distance Hydrogen Transmission

As global economies work to meet net-zero emissions targets, hydrogen has the capacity to bridge the gap between clean energy production and industrial power use by converting green power into fuel. Recent research has focused on optimizing the safety, cost, and energy efficiency using liquid hydrogen carriers (LHCs), such as ammonia, to facilitate the hydrogen supply chain. These studies investigate the cost of different carriers and transmission methods for specific supply scenarios. However, many studies leave out key cost variables along the supply chain, affecting the accuracy and adaptability of their results. The current study aims to provide a comprehensive techno-economic assessment of hydrogen transmission from Texas to New York via ammonia pipelines. This supply chain framework will fit into a larger model that compares levelized hydrogen transmission costs using different LHCs. Costs for hydrogen and ammonia feedstock, an ammonia synthesis plant, via the Haber-Bosch process, and a decomposition plant, via cracking were scaled from previous research. Due to ammonia's high corrosive properties and low boiling point (-33oC), costs to model a high-grade carbon steel pipeline with protective coatings and pump necessities were calculated to ensure safety and avoid significant pressure drop and flow issues. Initial results suggest levelized transmission cost of 14.52 $/kg-H2 at 50 tpd H2 throughput, driven by reactor plant CAPEX and operation costs. Additional sensitivity analysis shows that these costs drop to 8.77 $/kg-H2 at 500 tpd H2, in line with current ammonia research and closely comparable to other LHCs. The transmission energy intensity remains less than 1.0 MWth/e/MWH2, indicating net positive energy throughput despite the high energy intensity of ammonia synthesis and decomposition processes. Cost efficiency at higher hydrogen throughput indicates that ammonia could be an economically viable carrier for industry decarbonization.

Techno-Economic Assessment of Ammonia as a Liquid Carrier for Long Distance Hydrogen Transmission

NYU TANDON SCHOOL OF ENGINEERING

Aaryan Vira, Diya Sharma (PhD. Student Supervisor), Dharik Mallapragada (Principal Investigator)

Rationale

- Hydrogen transport has potential to connect clean energy generation and industrial power consumption (33.33 kWh/kg-H_2)
- Issues with transporting raw H_2:
  - Low density
  - Requires significant pressure for liquid storage/transport
  - Volatile and flammable
- Liquid hydrogen carriers → liquid at room temp (or under minimal pressure), less volatile, sustainable
  - **Ammonia,** MCH/Toluene, DBT, Ethanol/Acetic Acid

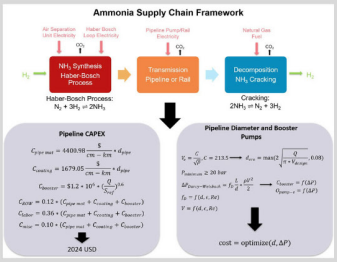
Ammonia as a Liquid Carrier

Advantages	Challenges
• High Gravimetric Density: 17.8 wt% H_2 • Liquid NH_3: 108.4 kg H_2/m^3 • Liquid H_2: ~80 kg/m^3 • Minimal cycle loss • Significant existing production and transmission infrastructure • One-way transmission	• Not Liquid at STP • -33°C or 8 bar saturation point • Highly corrosive • Synthesis and Cracking processes are energy intensive • Less energy efficient • More potential for emissions

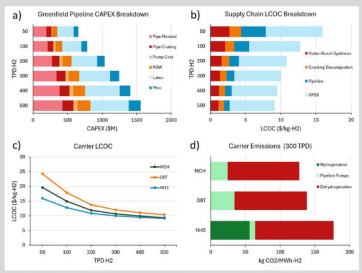
Objective

Design and assess NH_3 supply chain for hydrogen transmission that overcomes handling challenges to better understand NH_3 role as hydrogen carrier

Methods



TEA and Sensitivity Results



a) breakdown of the CAPEX of building an ammonia pipeline by cost category at increasing throughputs; **b)** breakdown of the levelized cost of carrier (LCOC) hydrogen delivery based on CAPEX and OPEX at increasing throughputs; **c)** LCOC comparison of different hydrogen carriers at increasing throughputs; **d)** comparison of carrier CO_2 emissions at 300 TPD throughput

Takeaways

- Pipeline CAPEX and OPEX are major drivers of hydrogen delivery costs → **optimization**
- NH_3 can be used as a cost-efficient hydrogen carrier
  - 500 TPD H_2 → **9.11 $/kg-$H_2$** (273.19 $/$MWh_{H2}$); Driven by OPEX
- Cost and emissions competitive with other hydrogen carriers

Next Steps

- Use NREL Renewable Energy Supply Curves database to spatially map wind and solar generation potential in Texas
- Use clean energy supply data to optimize cost of H_2 production via electrolysis

Selected References

Papadias, D. D., Peng, J.-K., & Ahluwalia, R. K. (2021). Hydrogen carriers: Production, transmission, decomposition, and storage. *International Journal of Hydrogen Energy*, *46*(47), 24169–24189. https://doi.org/10.1016/j.ijhydene.2021.05.002

Peecock, A., Hull-Bailey, B., Hastings, A., Martinez-Felipe, A., & Wilcox, L. B. (2024). Techno-economic assessment of liquid carrier methods for intercontinental shipping of hydrogen: A case study. *International Journal of Hydrogen Energy*, *94*, 971–983. https://doi.org/10.1016/j.ijhydene.2024.11.205

Poudel, S., Tanveer, S., Fang, Z., Lee, K., Lu, Y. R., Reddi, K., Sun, P., & Elgowainy, A. (2026). Techno-economic analysis of ammonia as hydrogen carrier. *International Journal of Hydrogen Energy*, *249*, 155975. https://doi.org/10.1016/j.ijhydene.2026.155975

Acknowledgements

I would like to thank Dr. Dharik Mallapragada and Diya Sharma for their invaluable guidance and mentorship. I would also like to thank the NYU Tandon Undergraduate Summer Research Program and the Visiting Scholar's Fellowship for making this research possible.

Industrial, Urban & Environmental Sustainability



Cleatus Oakes
Chemical and Biomolecular Engineering, B.S.
Spring 2029
California Institute of Technology

LTA Zeolite Films as Ion-Selective Membranes for Sustainable Lithium Extraction

Lithium is essential to modern batteries, yet the dominant solar-evaporation method for extracting it from natural brines is increasingly viewed as environmentally and operationally unsustainable due to heavy chemical use and substantial waste generation. This motivates the development of more selective, energy-efficient separation strategies. One promising approach employs LTA zeolites as molecular-sieving materials. LTA zeolites are aluminosilicate crystals with a well-defined cubic framework and characteristic pore apertures that can be synthesized as thin, continuous films on metal substrates. Building on this concept, we investigate whether LTA zeolite coatings can be grown directly on metallic surfaces and subsequently evaluated for their ability to selectively exchange Li^+ over competing ions. The LTA film was grown on a nickel substrate via hydrothermal crystallization over varying time periods. A precursor solution with the molar ratio $SiO_2 : Al_2O_3 : Na_2O : H_2O = 2.5 : 1 : 50 : 500$ was prepared, and the nickel support was roughened, cleaned, and submerged to initiate growth. Resulting films were assessed for thickness, adhesion, and continuity. Ion-exchange experiments were then conducted using Li^+/Mg^{2+} control solutions to quantify lithium uptake and evaluate competitive ion exclusion. Preliminary ion-exchange results demonstrate that LTA zeolite possesses a high decationization capacity, which enhances its lithium ion-exchange performance. These findings suggest that LTA zeolite films can serve as effective ion-selective membranes for lithium separation, supporting their potential as a scalable, environmentally sustainable alternative to conventional evaporation-based brine processing. Further optimization of film thickness and crystallization conditions may improve selectivity and recovery, advancing the development of more efficient lithium extraction technologies.

NYU TANDON SCHOOL OF ENGINEERING

Selective Lithium Uptake in LTA Zeolite Films via Ion Exchange

Cleatus Oakes[1]

1. UGSRP Visiting Student Fellow, California Institute of Technology

Abstract

LTA zeolite films are evaluated as selective ion-exchange materials for sustainable lithium recovery from brines. Conventional solar-evaporation extraction is increasingly considered environmentally and operationally unsustainable, motivating the search for more efficient separation methods. LTA zeolites, aluminosilicate crystals with uniform pore apertures and a cubic framework, can be synthesized as thin, continuous films on metal substrates and offer intrinsic size- and hydration-based ion selectivity. In this work, LTA films were hydrothermally grown on nickel supports and characterized for thickness, continuity, and adhesion. Ion-exchange experiments using Li^+/Mg^{2+} control solutions quantified lithium uptake and competitive ion exclusion. Preliminary results indicate that LTA zeolite exhibits a high decationization capacity and lithium selectivity, highlighting its potential as a scalable membrane material for environmentally sustainable lithium extraction.

Introduction

Lithium's low mass, high electrochemical potential, and ability to support high-energy-density storage make it central to modern batteries and the expansion of clean energy technologies [1]. However, the dominant solar-evaporation method for lithium recovery is increasingly viewed as environmentally intensive and operationally unsustainable. Because natural brines contain low lithium concentrations and high Mg^{2+}/Ca^{2+} levels, the process requires substantial chemical treatment, generating significant solid and liquid waste [2]. These challenges highlight the need for more selective, energy-efficient, and environmentally sustainable alternatives to evaporation-based lithium extraction.

One promising approach involves the use of zeolite-based ion-exchange materials, particularly those with the LTA framework [3, 4]. LTA-type zeolites are microporous aluminosilicates characterized by a well-defined cubic structure (Fig. 1A). In their sodium form (NaA), these zeolites possess a characteristic pore size (Fig. 1B) with a high ion exchange capacity, all without disrupting the integrity of the underlying aluminosilicate framework [5, 6]. Building on this foundation, we investigate whether similar zeolite materials can selectively ion exchange Na^+ with Li^+ over Mg^{2+}, thus creating a more sustainable method for lithium extraction.

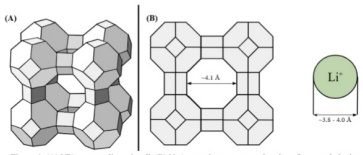


Figure 1. (A) LTA-type zeolite unit cell. (B) NaA pore size as compared to that of a strongly hydrated Li^+ ion. The pore can be selectively contracted or expanded through ion-exchange with K^+, shrinking to 2.9 Å, or Ca^{2+}, increasing to 4.4 Å [5].

Methods

The overall aim is to evaluate whether LTA zeolite coatings can enable selective ion exchange of Li^+ over Mg^{2+} in brine-like environments, thereby providing a more efficient pathway for lithium recovery.

Zeolite Layer Synthesis

- Nickel supports (0.15 mm thick) were cut into 44 × 44 mm squares, polished with 7 µm diamond paste, sonicated in acetone for 15 minutes, rinsed with DI water, air-dried, and the backside was covered with Teflon tape
- A synthesis solution with the composition $SiO_2 : Al_2O_3 : Na_2O : H_2O = 2.5 : 1 : 50 : 500$ was prepared by dissolving 17.51 g NaOH and 0.73 g $NaAlO_2$ in 80.4 g H_2O, followed by 1.34 g SiO_2 under stirring
- The nickel support was suspended face-down in the synthesis solution to ensure that only crystals with sufficient adhesion remained attached, after which 1 mL of a 5% v/v hexamethylenediamine dihydrochloride solution was added.
- Samples underwent hydrothermal crystallization at 50 °C, with growth monitored in 3-hour increments to characterize the development of the zeolite layer on the nickel baseplate

Ion Exchange Analysis

- The zeolite layer was positioned face-down in a solution of DI water while continuously stirred over a period of 6 hours to measure decationization capacity
- Brine-like conditions were then evaluated using the same assembly and a decationized zeolite layer, first with 150 mL DI water containing 60 mg/L of LiCl then with 150 mL DI water containing 60 mg/L of LiCl and 60 mg/L of $MgCl_2$
- Hourly samples were taken to quantify Na^+, Li^+, and Mg^{2+} concentrations and evaluate the progression of ion exchange

Results

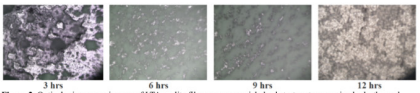


Figure 2. Optical microscopy images of LTA zeolite films grown on nickel substrates at successive hydrothermal crystallization periods. It highlights the progressive development of the zeolite film, showing how the surface morphology becomes increasingly textured and continuous as crystallization proceeds.

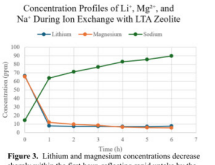


Figure 3. Lithium and magnesium concentrations decrease sharply within the first hour, reflecting rapid uptake by the zeolite, while sodium concentration increases steadily over time as Na^+ is released from the framework and exchanged with incoming Li^+ and Mg^{2+} ions.

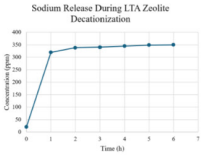


Figure 4. Sodium concentration in solution increases rapidly within the first hour, indicating fast initial Na^+ release from the zeolite framework.

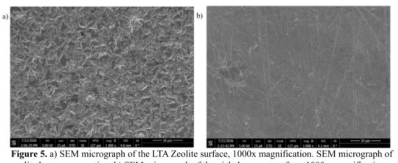


Figure 5. a) SEM micrograph of the LTA Zeolite surface, 1000x magnification. SEM micrograph of zeolite layer cross section. b) SEM micrograph of the nickel support surface, 1000x magnification.

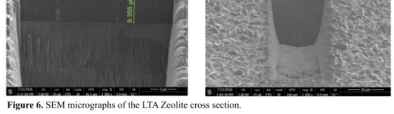


Figure 6. SEM micrographs of the LTA Zeolite cross section.

Discussion and Conclusion

The 12-hour hydrothermal crystallization period was selected because it produced the most uniform, continuous, and adherent LTA zeolite layer on the nickel substrate. Earlier growth intervals (3–9 hours) showed partial coverage and discontinuous crystal fragments, whereas the 12-hour film produced a stable surface morphology, thus ensuring reliable ion-exchange performance. This structural integrity was essential for the subsequent decationization and brine-like ion-exchange experiments. During decationization, sodium concentration increased sixteen-fold, demonstrating the high density of exchangeable Na^+ sites within the LTA framework and confirming that the synthesized film possessed a substantial cation-exchange capacity.

Under brine-like conditions, lithium and magnesium concentrations decreased sharply, by 89.5% and 91.5% respectively, within the first hour. Although both ions were taken up efficiently, the data show that lithium exchanges more rapidly than magnesium, consistent with its smaller hydrated radius and more favorable fit within the LTA pore structure. However, the zeolite was less selective than ideal: magnesium continued to exchange even after lithium levels had already dropped to near-equilibrium values. This suggests that future experiments should shorten the ion exchange duration to capture the window in which lithium uptake is dominant while magnesium exchange remains minimal. Optimizing contact time may enhance selectivity and further improve the application of LTA films as selective membranes for sustainable lithium recovery.

References

[1] Ngoy, K. R.; Lukong, V. T.; Yoro, K. O.; Makarebo, J. B.; Chukwuati, N. C.; Ibegbulam, C.; Etarigbo-Ikelegbe, O.; Ukoha, K.; Joe, T.-C. Lithium-Ion Batteries and the Future of Sustainable Energy: A Comprehensive Review. Renewable and Sustainable Energy Reviews 2025, 223, 115971. DOI: 10.1016/j.rser.2025.115971.

[2] Liu, Y.; Ma, B.; Lü, Y.; Wang, C.; Chen, Y. A Review of Lithium Extraction from Natural Resources. International Journal of Minerals, Metallurgy and Materials 2022, 30 (2), 209–224. DOI: 10.1007/s12613-022-2544-y.

[3] Kresge, C. T., Dhingra, S. S. Molecular Sieves. Kirk-Othmer Encyclopedia of Chemical Technology, 2004. DOI: 10.1002/0471238961.1315120511210912.a01.pub2.

[4] Wakeriaa, A. Recent Advances in Electrochemistry with Zeolitic Materials. Current Opinion in Electrochemistry 2020, 56, 101808. DOI: 10.1016/j.coelec.2025.101808.

[5] Artsishevski, N. A.; Grachev, A. L.; Kolesnyari, B. A.; Fedosov, D. A. Pervaporation of Water–Alcohol Mixtures on Cation-Exchanged LTA Zeolite Membranes. Petroleum Chemistry 2019, 59(8), 888–898. DOI: 10.1134/s0965544119080024.

[6] Campanelle, A., Liguori, B., Ferone, C., Capasso, D., Aprea, P. Zeolite-Based Monoliths for Water Softening by Ion Exchange/Precipitation Process. Scientific Reports 2022, 12 (1). DOI: 10.1038/s41598-022-07679-2.

Acknowledgements: I would like to thank my mentor, Professor Pavel Kots, for his continued guidance and support throughout the length of my project, as well as the Undergraduate Summer Research Program for its generosity in making this project possible.

Industrial, Urban & Environmental Sustainability



Daksh Mehta
Electrical Engineering and Computer Engineering (dual major), B.S.
Spring 2027
NYU Tandon

Citywide Embodied Carbon Benchmarking

Buildings are a major source of greenhouse gas emissions, and as operations become more efficient, the carbon committed at the moment of construction, known as embodied carbon, grows in importance. New York City is developing a citywide embodied carbon baseline that agencies can use to set and update benchmarks for new construction. Creating that baseline requires whole-building life cycle assessment data, but the real datasets available today are small and unevenly distributed across building types, heights, and structural systems, which makes benchmark values unstable. Synthetic datasets built from RSMeans building models cover the building stock more evenly but rest on standardized assumptions, and their absolute carbon values are biased. This project proposes a pathway where real office buildings from the Carbon Leadership Forum (CLF) benchmark dataset serve as real-world anchors, and RSMeans is used to estimate how their embodied carbon would change under controlled geometry changes. The key assumption is that even with this bias, RSMeans reliably captures how building geometry drives global warming potential (GWP) for the foundation, superstructure, and enclosure. Gradient boosting models trained on RSMeans learn how floor area changes shift element-level GWP, and these responses generate counterfactuals for each CLF building at sizes the real dataset does not cover, improving the balance of the benchmarking dataset. The models are validated by holding out entire wall types, which tests whether the learned responses transfer to configurations never seen in training, and prediction error stays below four percent for all three elements. Testing this by holding out each real building in turn, the range of values covered by the augmented benchmark contains the held-out building's true value more consistently than the range built from the raw sample alone. The outcome is a benchmark that reflects the full range of building sizes rather than the uneven coverage of the real dataset.

NYU TANDON SCHOOL OF ENGINEERING

BI LAB

Synthetic Augmented Benchmarking of Embodied Carbon

Daksh Mehta[1], Heng Quan[1], Semiha Ergan[1]

1. Tandon School of Engineering, New York University

Abstract

This study explores whether a small set of real buildings and a large set of synthetic ones can be combined to produce stable embodied carbon benchmarks. Real whole-building life cycle assessments are too few and too unevenly distributed to support reliable percentile values, while synthetic RSMeans building models cover the stock evenly but carry biased absolute carbon. Real buildings give the true carbon values, and a gradient boosting model trained on RSMeans concrete office buildings gives how much those values change when floor area changes. When tested on cladding types it was never trained on, the model predicts foundation, superstructure, and enclosure carbon within four percent error. Using the model this way on real buildings gives benchmarks that more often include the true value of a building left out of the data. The findings suggest a way to build benchmarks that span the full range of building sizes and grow more accurate as new projects are collected.

Introduction

A benchmark tells a building design team what carbon intensity is normal for their building type early in design, when the structure and cladding that drive embodied carbon can still be changed [1]. Broken out by element, it also shows where to cut carbon emissions specifically. However, these values are percentiles drawn from whatever buildings have been measured, and the largest North American dataset, the Carbon Leadership Forum's WBLCA benchmark study, holds fewer than 300 buildings [2], which leaves very few projects once split by the building type and structural system a benchmark needs.

Methodology

To close this gap, synthetic data is used to support real data. How carbon intensity responds to floor area is learned from the 702 synthetic RSMeans [3] concrete offices alone. Buildings sharing height class, cladding system, story count, story height, and perimeter ratio are matched into 1,404 pairs that differ only in floor area. For each of the three elements, foundation, superstructure, and enclosure, a gradient boosting model (GBM) is fitted to predict the change in carbon intensity as a log ratio from four inputs: the ratio of new to original floor area, the original floor area, the height class, and the building's own starting carbon intensity. One GBM is fitted per element, using the same hyperparameters throughout, since the three elements respond to size differently.

Validation holds out one cladding system at a time, trains on the rest, and tests on the one left out, repeating for all four cladding systems. Since the model never trains on that cladding, this checks whether it generalizes to cladding it has not seen, rather than fitting the training data.

Once validated, the final models are trained on all pairs and generate the new data points. They are used on the 12 real concrete offices, predicting each one's carbon intensity at other floor areas typical for its height class, producing 36 new data points. Internally, the model predicts a log ratio, which is converted using exp to find the actual ratio, then multiplied by the building's real value to get the predicted carbon intensity. Figure 1 summarizes this pipeline.

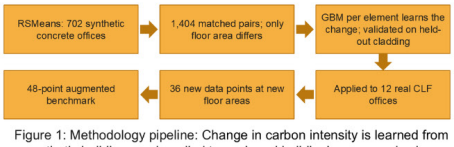


Figure 1: Methodology pipeline: Change in carbon intensity is learned from synthetic buildings and applied to each real building's measured value.

Results

Tested on cladding held out of training, the model's error, measured as mean absolute percentage error (MAPE) on the predicted ratio, stays under four percent for all three elements: 0.7 for foundation, 3.9 for superstructure, and 3.8 for enclosure.

Because the RSMeans grid is complete, its true median and 80th percentile carbon intensity are known for each element. Drawing 300 size-imbalanced subsamples from it shows that adding counterfactuals moves the estimated median and 80th percentile closer to these true values.

On real data, leaving out each of the 12 buildings in turn, the 10th to 90th percentile band built from real data alone contains the held-out building's actual measured value in 8 of 12 cases (66.7%) for every element. Adding counterfactuals raises this to 9 of 12 (75%) for superstructure and 10 of 12 (83.3%) for enclosure; foundation stays at 8 of 12 (66.7%).

Limitations

The model only captures how carbon intensity responds to floor area, not to other design choices that also drive embodied carbon. Validation so far holds out only the cladding system, a more complete test would also hold out other variables such as story count. Lastly, the real sample used is just twelve concrete office buildings, so more testing would be needed.

Works Cited

[1] Röck, M., et al. (2020). Embodied GHG emissions of buildings: The hidden challenge for effective climate change mitigation. Applied Energy, 258, 114107.
[2] Benke, B., et al. (2025). A harmonized dataset of high-resolution embodied life cycle assessment results for buildings in North America. Scientific Data, 12.
[3] Gordian. (2025). RSMeans Data Online. Gordian.

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.

Industrial, Urban & Environmental Sustainability



Daniel Zheng
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon

On The Need For Benchmarking Space-Charge-Limited Current Measurements for Metal-Halide Perovskites

Space-charge-limited current (SCLC) measurement is one of the most widely used steady-state techniques for estimating charge-carrier mobilities and electronic trap densities in semiconductors, including metal-halide perovskites (MHPs) – a leading material for next-generation photovoltaics and optoelectronics. At higher voltages, SCLC measurement-based current density–voltage (J-V) curves follow a square-law dependence that can be fit using the Mott-Gurney (MG) law to extract these parameters. The MG law depends on multiple assumptions, most importantly that it applies only to single-carrier devices. Yet despite the apparent simplicity in measurement, electric field-induced mobile ion migration in MHPs complicates the reliability and accuracy of SCLC-derived parameters. Here, we analytically model continuous SCLC measurements on $MAPbI_3$ devices, a well-characterized I^- ion conductor, to demonstrate that the presence of mobile ions significantly distorts the J-V characteristics under dynamic measurement conditions – specifically scan rates, applied electric field, and bias looping. Repeated electric field cycling progressively lowers the extracted charge carrier mobility, which we hypothesize arises from a reduction in current density due to mobile-ion-induced internal field screening. Interestingly, this progressive drop depends only on the time spent under bias and is independent of scan rate or the maximum applied electric field. We also highlight that electric field looping induces a recoverable metastable state and that thermal annealing can restore near initial current density values and subsequently mobility, implying that the accumulated mobile ions redistribute in the bulk of the MHPs. This work reveals that an MHP's measurement history can dynamically affect parameters extracted from simplistic SCLC analysis, reflecting known but poorly understood mixed electronic-ionic contributions. This highlights the need for measurement standards that account for mobile-ion dynamics to ensure SCLC-derived transport and defect parameters are physically meaningful and reproducible.

NYU TANDON SCHOOL OF ENGINEERING

Is Space-Charge-Limited Current a suitable measurement standard for mixed-ionic electronic semiconductors?

Daniel Zheng[1,2], Souhardya Bera[2], Seth McPherson[2], André D. Taylor[2], Ayaskanta Sahu[2], Jason A. Röhr[1]

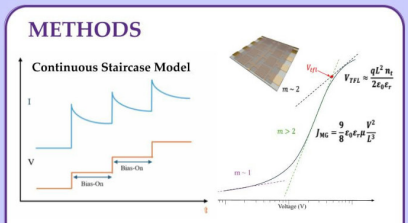
[1]General Engineering, Tandon School of Engineering, New York University
[2]Department of Chemical and Biomolecular Engineering, Tandon School of Engineering, New York University

INTRODUCTION

Space-Charge-Limited Current (SCLC) is a fundamental steady-state measurement technique to estimate transport and defect properties in emerging semiconductors – including metal halide perovskites (MHP's) and quantum dots (QD's) – leading materials in next-generation photovoltaics and optoelectronics.

- Contribution of electronic and ionic species influence the measured current in both MHP's & QD's during biasing.
- The Mott-Gurney law is often used to extract mobility on the quadratic regime of the current density (J)-voltage (V) curve.
- However, there still remain several issues that raise concerns about the reliability of reported values.

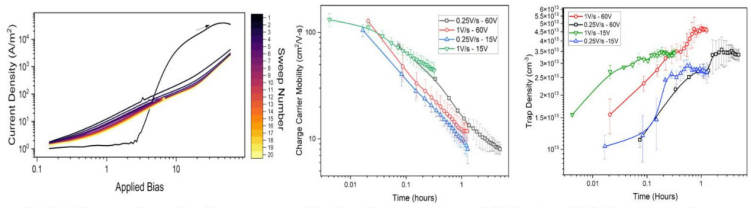
METHODS



Scan-rate variable measurements with stepwise increases allow for significant ion contribution and space-charge buildup leading to the origin of hysteresis.

- Mobility, μ, is extracted from fitting the Mott Gurney Law in the regime where slope ~ 2
- Trap density is extracted from estimated VTFL at the point of intersection between the region of slope > 2 & slope ~ 2.

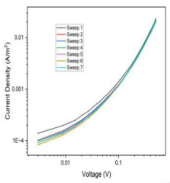
METHYLAMMONIUM LEAD IODIDE RESULTS



- Positive Sweeps shows the decrease of Current Density with Sweep Number (Time)
- Shift in V_{TFL} indicative of a changing trap density
- Varying the scan rate and highest applied bias by 4x shows no measurable difference in charge carrier mobility (μ) reduction
- Mobility reduction spans nearly two orders of magnitude over a time-scale of hours.
- The fall in mobility can be attributed to mobile ion-induced internal field screening.

PbS QD CASE STUDY

- Taking inspiration from MHP's, the same SCLC measurements were performed on PbS colloidal quantum dot devices.
- Extracted mobility values drop slightly from 3.1 × 10⁻⁸ cm²/(V·s) to 2.8 × 10⁻⁸ cm²/(V·s) after the first sweep indicating no dependence on measurement history.



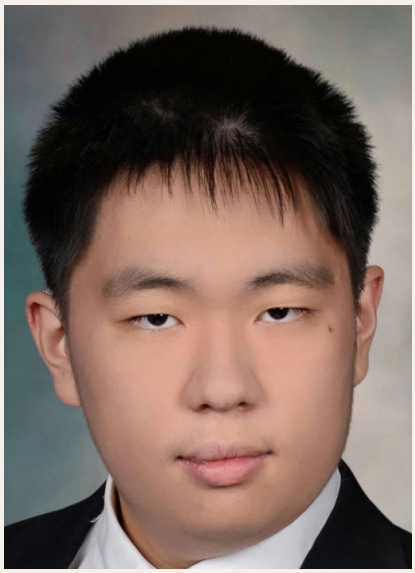
CONCLUSIONS

- Mobile ion migration under bias complicates the analysis of measured current density
- Charge-carrier mobility (μ) is only dependent on the measurement history of the device for MHPs
- Key parameters regarding optoelectronic properties remain extractable from QD's, despite fundamentally different transport mechanisms from traditional semiconductor materials used in SCLC.

REF

Special thanks to the UGSRP (Undergraduate Summer Research Program) and Hybrid Nanomaterials Laboratory for providing this opportunity.

Industrial, Urban & Environmental Sustainability

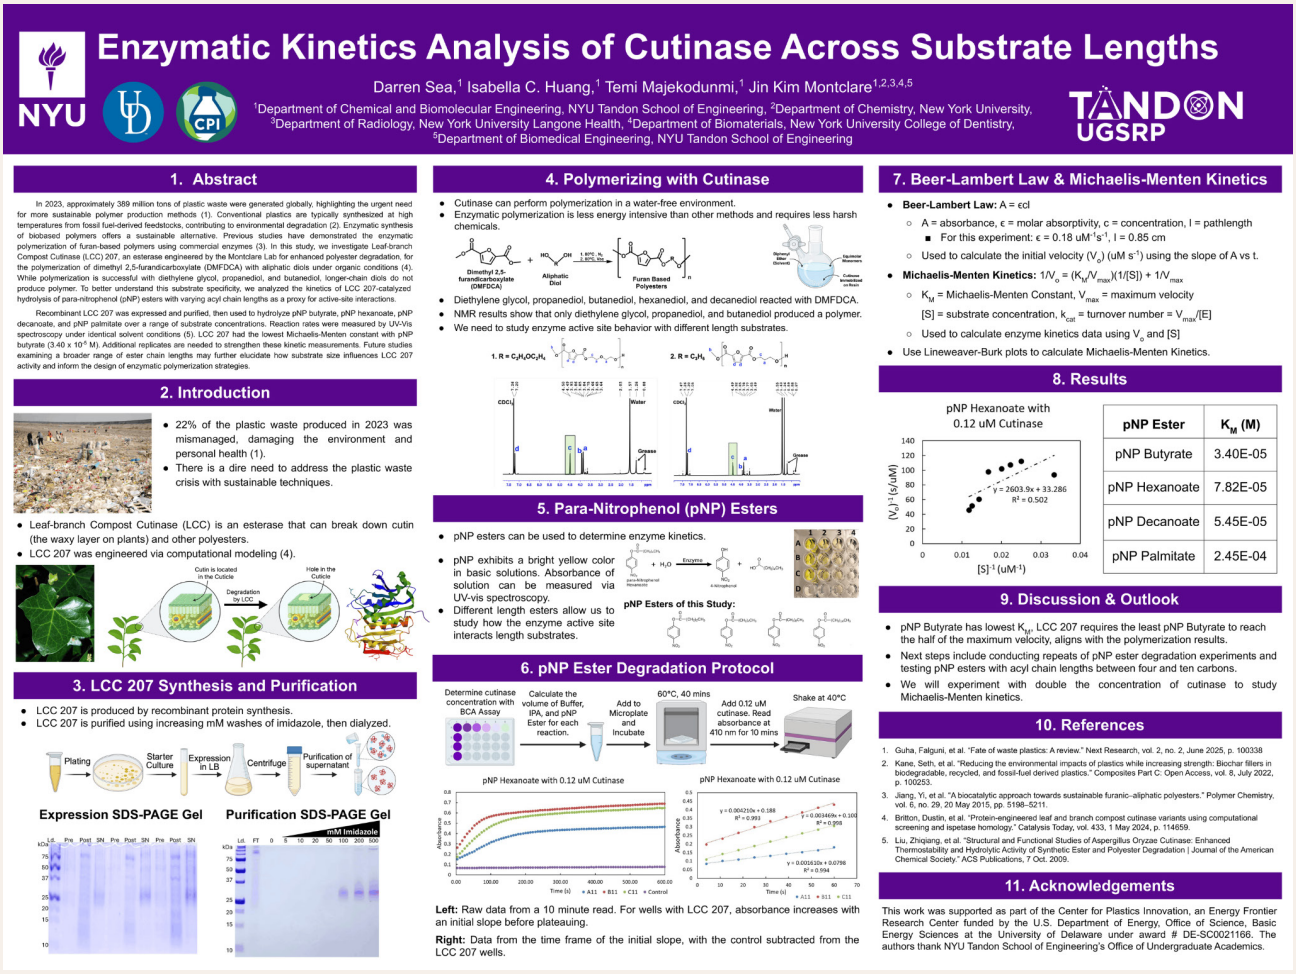


Darren Sea
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon

Enzymatic Kinetics Analysis of Cutinase Across Substrate Lengths

In 2023, approximately 389 million tons of plastic waste were generated globally, highlighting the urgent need for more sustainable polymer production methods (1). Conventional plastics are typically synthesized at high temperatures from fossil fuel-derived feedstocks, contributing to environmental degradation (2). Enzymatic synthesis of biobased polymers offers a sustainable alternative. Previous studies have demonstrated the enzymatic polymerization of furan-based polymers using commercial enzymes (3). In this study, we investigate Leaf-branch Compost Cutinase (LCC) 207, an esterase engineered by the Montclare Lab for enhanced polyester degradation, for the polymerization of dimethyl 2,5-furandicarboxylate (DMFDCA) with aliphatic diols under organic conditions (4). While polymerization is successful with propanediol and butanediol, longer-chain diols do not produce polymer. To better understand this substrate specificity, we analyzed the kinetics of LCC 207-catalyzed hydrolysis of para-nitrophenol (pNP) esters with varying acyl chain lengths as a proxy for active-site interactions.

Recombinant LCC 207 was expressed and purified, then used to hydrolyze pNP butyrate, pNP hexanoate, pNP decanoate, and pNP palmitate over a range of substrate concentrations. Reaction rates were measured by UV-Vis spectroscopy under identical solvent conditions (5). LCC 207 exhibited the highest catalytic efficiency toward pNP hexanoate ($3.2 \times 10^3\ M^{-1}s^{-1}$), followed by pNP decanoate ($2.4 \times 10^3\ M^{-1}s^{-1}$), pNP butyrate ($1.6 \times 10^3\ M^{-1}s^{-1}$), and pNP palmitate ($6.2 \times 10^2\ M^{-1}s^{-1}$). Additional replicates are needed to strengthen these kinetic measurements. Future studies examining a broader range of ester chain lengths may further elucidate how substrate size influences LCC 207 activity and inform the design of enzymatic polymerization strategies.



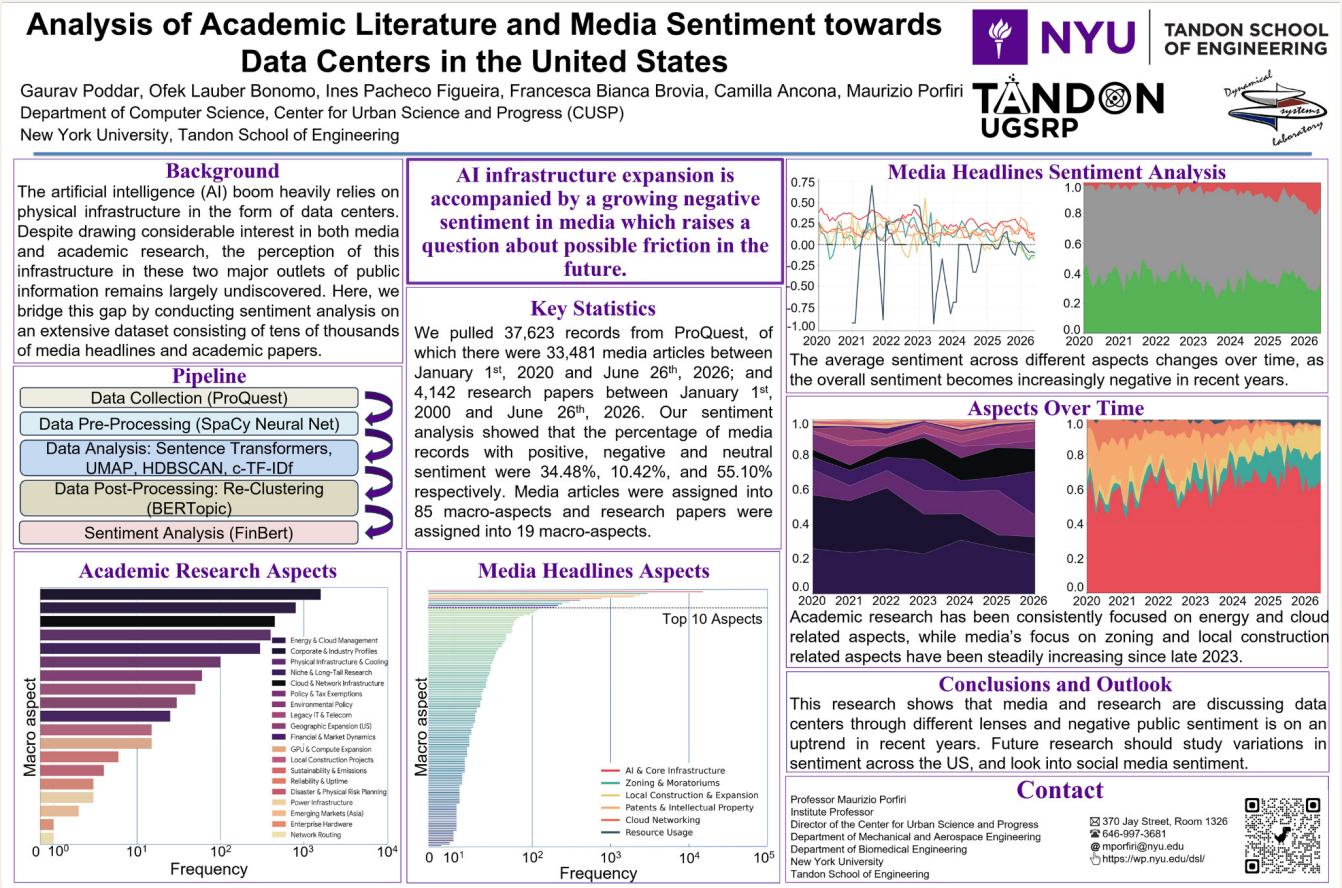
Industrial, Urban & Environmental Sustainability



Gaurav Poddar
Computer Science, B.S.
Spring 2028
NYU Tandon

Tracing the AI Expansion: A Multidimensional Analysis of Academic Literature, Media Sentiment, Data Center Growth, and Public Search Behavior

The explosive growth of artificial intelligence depends heavily on physical data center infrastructure, yet the interplay between technical innovation, academic research, public media sentiment, and search behavior remains poorly understood. This study quantifies the lifecycle of AI expansion as abstract compute progress translates into physical infrastructure, local regulatory friction, and public information seeking. I synthesized a 26-year multimodal dataset spanning academic literature, industry news, and public search trends. Utilizing financial domain Large Language Models (FinBERT) and dense sentence embeddings, we measured sentiment polarity and semantic drift. To model temporal and spatial dynamics, we integrated Cross-Correlation Functions, Granger Causality testing, Discrete-Time Markov Chains, and empirical Event Studies across geographic regions. Academic compute research statistically leads industry media coverage by 3 months ($p = 0.0001$), with GPU literature undergoing a massive semantic shift post-2020 (cosine distance $= 0.86$). In news media, AI infrastructure enthusiasm acts as a leading indicator for local zoning and resource backlash up to 11 months downstream. Negative media outrage decays on a predictable exponential curve with a half-life of 30.7 days. Furthermore, spatial Markov modeling proves a sharp structural divergence (Frobenius norm $= 1.34$) between execution-driven Tier-1 hyperscale markets and policy-constrained emerging markets. AI infrastructure expansion follows a highly structured, predictable lifecycle spanning academic discovery, media coverage, and public interest. By quantifying the temporal lags, recovery rates, and regional bottlenecks of data center expansion, this framework enables researchers, policymakers, and industry leaders to anticipate physical constraints and public perception as digital infrastructure scales.



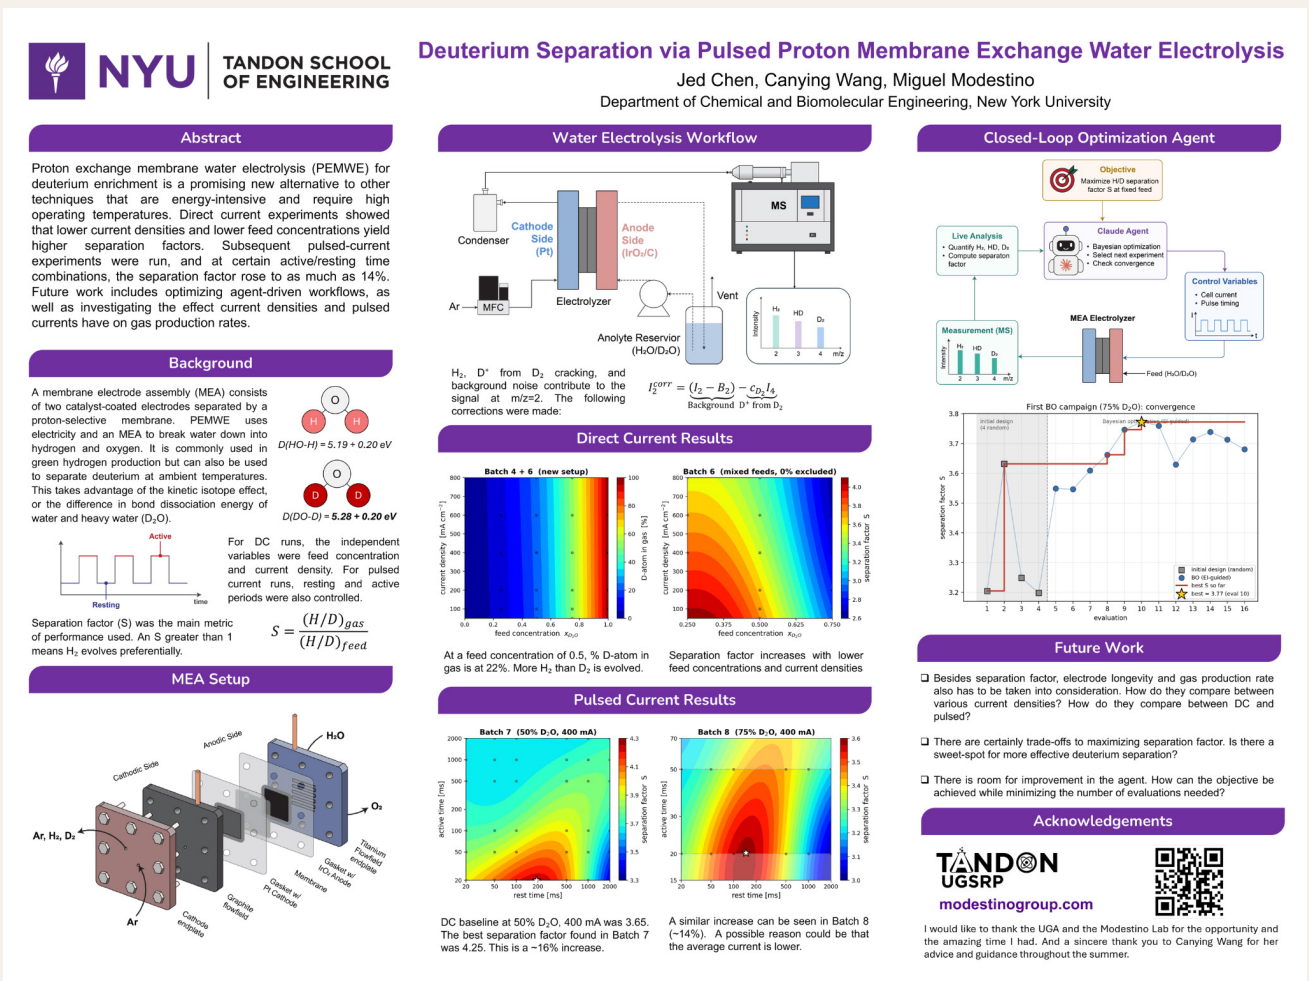
Industrial, Urban & Environmental Sustainability



Jed (Ethan) Chen
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon

Pulsed Proton Exchange Membrane Water Electrolysis for Deuterium Separation

Heavy water (D_2O) is produced in quantities of hundreds of metric tons annually for nuclear power applications, most commonly through isotopic exchange and distillation techniques that are energy-intensive and require extreme operating temperatures. This project explores proton exchange membrane water electrolysis (PEMWE) as a promising alternative: it already underpins green hydrogen production, operates at ambient temperature, and achieves effective isotopic separation. We hypothesized that a pulsed current would outperform direct current (DC) operation by allowing active sites to be replenished during the resting portion of each cycle. Using a membrane electrode assembly comprising Pt || IrO_2/C catalysts and a Nafion membrane, we first established a DC baseline: lower current densities and lower feed concentrations produced higher separation factors. A subsequent pulse-current parameter sweep showed that, at matched peak current density and feed concentration, certain active/resting time combinations raised the separation factor by approximately 14%. This improvement could be attributable to the lower time-averaged current rather than to pulsing itself—a distinction that will guide how we interpret future results. Additional experiments, assisted by an autonomous optimization agent, are underway to explore this parameter space more broadly, with hydrogen production rate and other metrics still needed before drawing definitive conclusions.



Industrial, Urban & Environmental Sustainability



Jed Liu
Mechanical Engineering, B.S.
Spring 2028
NYU Tandon



Reet Barua
Mechanical Engineering, B.S.
Spring 2028
NYU Tandon

Investigation of Materials Using Density Functional Theory

The Gibbs-Thomson effect describes how the phase equilibrium of a nanomaterial alters due to elevated surface energy. Specifically, this thermodynamic phenomenon leads to a significant melting point depression as a material's surface-area-to-volume ratio increases at the nanoscale. Melting point depression is highlighted by the behavior of pure gold, which drops from a bulk melting point of 1337 K to near 900 K for a nanoparticle diameter of just 30 Å. This research aims to simulate and model nanoscale melting point depression using the MedeA software on the High Performance Computing (HPC) facility. Supercells of varying dimensions were engineered in the platform to run molecular dynamics simulations for gold, aluminum, and copper nanostructures. The simulation data revealed a non-linear drop in melting temperature that closely aligned with external experimental results, validating the MedeA software as an accurate predictive framework for evaluating the thermal properties of metallic nanostructures. The software's atomic modeling capabilities can expand this research to practical engineering challenges, including metal 3D-printing optimization and airborne fume risk modeling.

NYU TANDON SCHOOL OF ENGINEERING

Investigating Melting Points of Gold, Aluminum and Copper Using Molecular Dynamics Simulations

Jed Liu, Reet Barua, Vedant Nadoda, Rakesh K. Behera
Department of Mechanical and Aerospace Engineering, NYU Tandon School of Engineering

Abstract

The Gibbs-Thomson effect describes how the phase equilibrium of a nanomaterial alters due to elevated surface energy, leading to significant melting point depression as nanoparticle size decreases. In ambient environments, metal nanoparticles naturally undergo surface oxidation, which alters their thermodynamic stability. Isolating the impact of oxidation requires first establishing a baseline melting point under oxygen-free conditions, which is quite challenging to perform experimentally. Hence, this research aims to simulate and model nanoscale melting point depression using the LAMMPS simulation engine integrated within the MedeA software platform. As an initial benchmark, two-phase molecular dynamics simulations were conducted to evaluate candidate interatomic potentials for bulk gold, aluminum, and copper. The simulation results predicted melting temperatures with relative errors of approximately 7% for aluminum and 15% for gold, validating the selected interatomic potentials as a reliable baseline for evaluating the thermal properties of metallic nanostructures.

Background

Nanoparticle vs. Bulk Behavior: In a bulk crystal lattice, interior atoms are highly stable since they are fully coordinated by their nearest neighbors. In contrast, surface atoms are less stable as they lack neighbors in certain directions. A majority of atoms in a bulk material lie at its fully coordinated core, while a significant fraction of atoms in a nanoparticle reside at its under-coordinated surface. Since the average coordination number decreases with increasing surface-to-volume ratio, the Gibbs-Thomson effect predicts a decrease in melting temperature as particle diameter decreases to the nanoscale.

Table 1: Aluminum, Gold, and Copper Experimental Bulk Melting Reference Values

	Experimental Bulk Melting Reference Values (K)
Aluminum	933.47
Gold	1337.33
Copper	1357.77

$$E_i = F_\alpha\left(\sum_{j\neq i}\rho_{\alpha\beta}(r_{ij})\right) + \frac{1}{2}\sum_{j\neq i}\phi_{\alpha\beta}(r_{ij})$$

Equation 1: Embedded Atom Method Formula

Methodology

All simulations were run on the **MedeA software**. MedeA allows for the atomistic and molecular modeling of materials, making an efficient and reliable way to simulate and test materials without dealing with imperfection found in real-life experiments.

To automate the phase coexistence simulation protocol, a multi-stage workflow (shown above) was constructed in MedeA using the **LAMMPS engine**.

1. The 4 x 4 x 20 supercell was generated under the conditions set in the MedeA flowchart.

2. A custom LAMMPS script is executed to selectively heat the supercell according to the preset constraints.

3. A while loop then executes until it can be evaluated that the crystal structure grew or melted – which can then be determined to be the thermodynamic melting point.

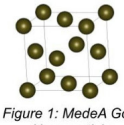


Figure 1: MedeA Gold Nanoparticle

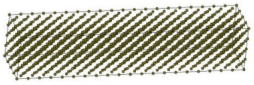


Figure 2: MedeA 4 x 4 x 20 Gold Supercell

Results

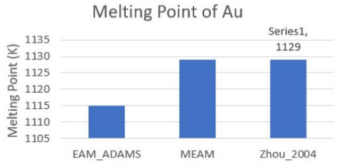


Figure 3: Bulk Aluminum Melting Point Predictions Across Interatomic Potentials

Gold (Best Potential): EAM_ADAMS, Zhou_2004
Calculated Percent Error: 15.58%

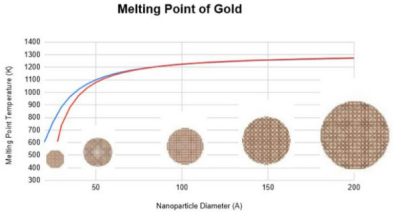


Figure 4: Melting Point of Gold Nanoparticle Graph Created From Literature Values

The above graph was generated after reviewing literature values investigating the melting point of gold at different nanoparticles.

The results compiled above in the bar chart show the experimental melting point on the y-axis, simulated by MedeA, and the mathematical algorithms (potentials) used in LAAMPS to carry out the simulation on the x-axis. The same graphs were generated for gold and copper. The graph of the literature values is also shown above, and is indicated by the red and blue lines (two different ways to calculate melting point temperature). Compared to the values obtained from literature review of previous studies, the calculated percent error is as shown above. The same chart was created and percent error was calculated for aluminum and copper nanoparticles.

Future Work

As researchers initially new to the High-Performance Computing (HPC) workflows, atomistic simulation methodologies, and advanced modeling packages, significant effort was dedicated to understanding these platforms. To examine the impact of oxidation and the Gibbs-Thomson effect, there is still substantial simulation work needed across surface oxygen concentrations and varying nanoparticle diameters. As of now, we focused primarily on gold due to its resistance to oxidation.

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project. The authors would also like to acknowledge the NYU HPC facility for computational resources.

Industrial, Urban & Environmental Sustainability



Hyunggeun Noah Yoon
Mechanical Engineering, B.S.
Spring 2029
NYU Tandon

Multiphase Coexistence in Complex Mixtures

Polydisperse systems have a continuous range of chemically distinct species characterized by distributions of factors including molecular weight, charge, or polarity, making direct thermodynamic analyses, such as phase equilibrium calculations, computationally impossible. To address this, we developed a Python-based computational structure that implements the ideas behind moment-free energies from Sollich, Warren, and Cates (2000), which finds that the excess free energy of many models, particularly mean-field frameworks, is "truncatable"– they depend only on a finite number of moments associated with the attribute's distribution. Free-energy-like functions are visualized as landscapes in which Monte Carlo simulations explored low-energy, high-probability regions and symbolic differentiation produced accurate curvature classification; specifically, unstable or saddle-like regions connected with phase separation were identified by negative Hessian eigenvalues. Tilting the initial landscape f(r) to g(r), where g(r)=f(r)-~f(r0)(r-r0), allowed us to study coexistence pathways from approximated reference points without changing local curvature; thus, the application of one-dimensional search paths originating from negative eigenvectors associated with those points to identify the location of minima (that serve as candidate coexisting phases). The found line-search minima were grouped by k-means clustering, enabling the visualization of distinct basins through contour maps, curvature-based cluster plots, and heat maps. The constructed framework applies to numerous industrially relevant systems, as it reveals likely phase basins and pathways for phase separation, paving the way for better prediction of how complex mixtures separate into coexisting phases and enabling applications in material design, including polymers and colloids.

NYU TANDON SCHOOL OF ENGINEERING

TANDON UGSRP

Phase Coexistence in Complex Mixtures

Noah Yoon[1], Phillip Rauscher[2]

1. Department of Mechanical and Aerospace Engineering 2. Department of Chemical and Biomolecular Engineering

Introduction

Polydisperse systems have a continuous range of chemically distinct species characterized by distributions of attributes, making direct thermodynamic analyses computationally impossible. To address this, we developed a Python-based computational structure based on the idea that excess free energy of many models depend only on a finite number of moments associated with the attribute's distribution. Tilting the initial landscape allowed us to study coexistence pathways from approximated reference points without changing local curvature; thus, the application of one-dimensional search paths originating from negative Hessian eigenvectors associated with those points to identify the location of minima (that serve as candidate coexisting phases). The found line-search minima were grouped by k-means clustering, enabling the visualization of distinct basins through several characteristic plots. The constructed framework applies to numerous industrially relevant systems, as it reveals likely phase basins and pathways for phase separation, paving the way for better prediction of how complex mixtures separate into coexisting phases.

Method

The moment free energy method by Warren, Sollich, and Cates (2000) [1] reduces a polydisperse system's thermodynamic descriptions to a finite set of statistical moments. Any two coexisting phases are connected by a region of instability, thus we searched for unstable areas and transformed them into saddle points, which led us to candidates for coexisting phases that this workflow aims to find. The primary function utilized throughout the duration of the project is Himmelblau's function:

$$f(x,y) = (x^2 + y - 11)^2 + (x + y^2 - 7)^2.$$

Monte Carlo simulations from the Python package *littlemcmc* explored low-energy, high-probability regions while symbolic differentiation produced accurate curvature classification. K-means clustering was done for points with negative Hessian eigenvalues; the centers of each cluster with one negative Hessian eigenvalue were used as approximated points of instability transformed to saddles that started the negative Hessian eigenvector line search for minima in the tilted landscapes. Given the original landscape, the tilted landscape is denoted as $g(r)$, where:

$$g(r) = f(r) - \nabla f(r_0) \cdot (r - r_0),$$

and the one-dimensional line search along the negative Hessian eigenvector directions is denoted by:

$$r(t) = r_0 + t v_0,$$

where t is a small scalar parameter and v_0 is the direction(s) of the eigenvector associated to the negative Hessian eigenvalue.

Result

Line searches from the four cluster centers were able to find two minima in their respective tilted landscapes, each of which are nearby the original function's known local minima. Doing further clustering for these eight points should provide accurate approximations for solving phase coexistence equations. Though not sufficient on its own, it is necessary for coexisting phases to have positive curvature, as shown in Figure 1.

Figure 1: Result of Negative Hessian Eigenvector Line Search for Minima from cluster centers of points with one negative Hessian eigenvalue

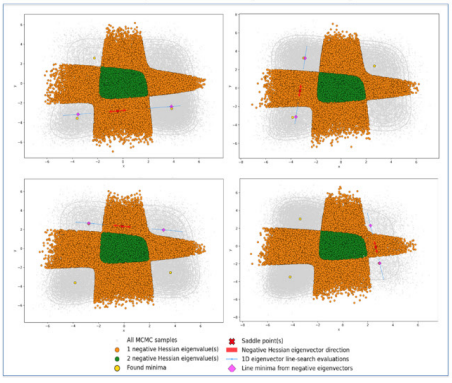
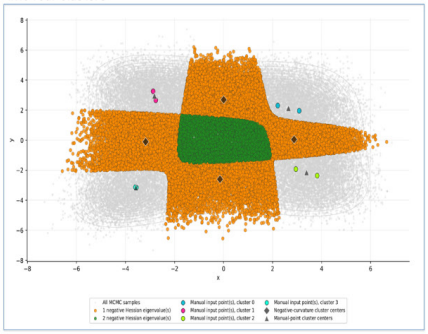


Figure 2: K-means clustering of all eight points found from line search into four clusters



Discussion

- Negative Hessian Eigenvalues determine unstable areas where phase separation could occur.
- Results can vary depending on presets like temperature, search bounds, number of clusters, and sampling size. For k-means clustering, run either elbow test or silhouette score for a particular function for most optimal k value.
- Graphical visualization may become more difficult for landscapes in higher dimensions.

Future Work

- Test other physically realistic free-energy models of polydisperse systems with the workflow.
- Continue developing visualizations for landscapes in higher dimensions.
- Integrate automatic selection of k value into workflow.

Works Cited

[1] P. Sollich, P. B. Warren, and M. E. Cates, *Moment Free Energies for Polydisperse Systems*, in Advances in Chemical Physics, vol. 116, I. Prigogine and S. A. Rice, Eds. Hoboken, NJ: John Wiley & Sons, 2001, pp. 265–336. https://doi.org/10.1002/9780470141762.ch4

Acknowledgements

A big thank-you to Professor Phillip Rauscher for his generous aid and mentorship all throughout the duration of the project and NYU Tandon School of Engineering's Office of Undergraduate Academics for providing funding for the project.

Industrial, Urban & Environmental Sustainability



Ian Yum
Electrical Engineering, B.S.
Spring 2028
NYU Tandon

True Colors: Color Rendering Index and Indoor Photovoltaics Efficiency

Photovoltaics are a promising technology for powering indoor ecosystems; however, indoor photovoltaics (IPVs) must be designed to match specific lighting sources and conditions. Previous studies have highlighted the importance of correlated color temperature (CCT) in determining the power efficiency and optimum bandgap energy of IPVs under illumination of white-light LEDs, however the role of the color rendering index (CRI) remains less understood. We investigated the impact of CRI and CCT on IPVs to inform future design choices by employing a computational approach to calculate maximum power-conversion efficiency and optimum bandgap values of IPVs when illuminated by different commercial white-light LEDs with varying CCT, CRI, and illuminance levels. Our results highlight that high-CRI LEDs require materials with lower bandgap energies for optimum efficiency as there is a shift towards red in the high CRI irradiance spectra: 1.75 to 1.80 eV as opposed to 1.76 to 1.90 eV in low-CRI environments. The performance of commercial solar cells was also evaluated, and we show that despite IPVs having optimal performance with wide-bandgap material under lower CRI lighting, mature technologies like silicon and CdTe benefit from high-CRI illumination. Additionally, we found that next-gen, wide-bandgap materials would offer higher efficiencies in typical indoor lighting scenarios. We conclude that both CRI and CCT need to be accounted for when developing IPVs because the broader irradiance spectra of high-CRI LEDs shift the optimum bandgap energies to lower values and reduce overall device efficiency. These findings can guide IPV deployment in environments like smart homes, retail, industrial monitoring, and other areas with varying lighting choices.

NYU

TANDON SCHOOL OF ENGINEERING

True Colors: Color Rendering Index and Indoor Photovoltaics Efficiency

Ian Yum, Aditi Sharma, Alex Guaman, Gavin Akmam, Jason A. Röhr

General Engineering, Tandon School of Engineering, New York University

Abstract

Indoor photovoltaics (IPVs) are a promising technology for powering indoor ecosystems. We investigated the impact of the color rendering index (CRI) and color correlated temperature (CCT) on IPVs by employing a computational approach to calculate maximum power-conversion efficiency and optimum bandgap values of IPVs when illuminated by different commercial white-light LEDs with varying CCT, CRI, and illuminance levels. We conclude that both CRI and CCT need to be accounted for when developing IPVs because the broader irradiance spectra of high-CRI LEDs shift the optimum bandgap energies to lower values and reduce overall device efficiency.

Background

The increasing number of IoT devices calls for a renewable power source, such as indoor photovoltaics (IPVs). White-light LEDs, commonly used in indoor environments, emit at lower intensities within a narrow spectral range. As a result, theoretical optimal bandgaps of indoor solar cells shift to ~1.8-2.0 eV (~1.1 eV optimal for crystalline silicon; relies on infrared absorption). Newer technologies can tune bandgaps much better for indoor emission profiles and have achieved efficiencies of 30-45%, which still falls short of the theoretical maximum. We Investigate how color temperature (CCT) and the color rendering index (CRI) affect theoretical efficiencies and optimum bandgaps for IPVs to understand how to optimize this technology.

Figure 1: Color Correlated Temperature

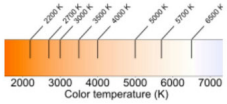


Figure 2: Color Rendering Index

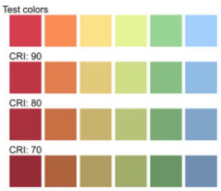


Figure 3: Photon Flux Density of the AM1.5G Spectrum and Bandgap Edge of Crystalline Silicon

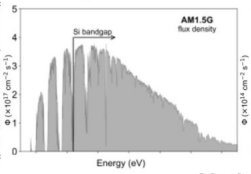
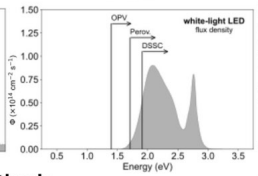


Figure 4: Photon Flux Density and Bandgap Values for DSSC, Perovskites, and OPV Materials



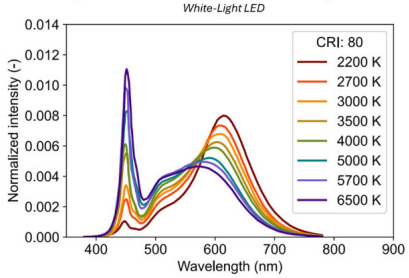
Methods

Commercial LED spectra with varying CCT and CRI were obtained. The spectra was then normalized to unit area (fig. 5) and scaled to a fixed illuminance (eq. 1) using the CIE 1931 photopic response. A detailed-balance model [1] was used to calculate theoretical maximum power conversion efficiency (eq. 2) of ideal solar cells under the scaled spectra.

$$a = \frac{E^{\text{target}}}{\kappa \int_{\lambda_1}^{\lambda_2} \varphi_{\text{norm}} V(\lambda)\, d\lambda} \qquad (1)$$

$$\eta = \frac{p_{\text{out}}}{p_{\text{in}}} \qquad (2)$$

Figure 5: Normalized Irradiance Spectra of Samsung LM301Z+ White-Light LED



Findings

We confirm that varying CCT impacts device performance, lower CCT values lead to higher efficiencies and lower optimal bandgaps. We found higher CRI LEDs require lower bandgap energies for optimum efficiency. IPVs preform best with wide-bandgap materials under low CRI, but mature technologies like silicon and CdTe would benefit from high-CRI illumination.

Figure 6: Theoretical Efficiency and Optimum Bandgaps at 500 lx

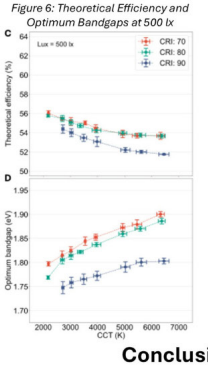
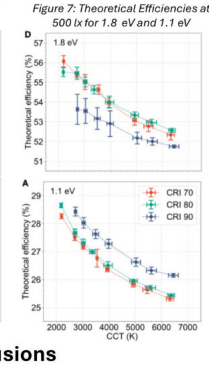


Figure 7: Theoretical Efficiencies at 500 lx for 1.8 eV and 1.1 eV



Conclusions

CRI and CCT must be accounted for in IPV development. Next-generation wide-bandgap materials offer better efficiencies indoors. Future work includes accounting for Shunt Resistance and Non-radiative recombination.

Works Cited

[1] Sharma, A., Guaman, A. A., & Röhr, J. A. (2026). *On the role of color temperature and color rendering index of white-light LEDs on the theoretical efficiency limit of indoor photovoltaics*. arXiv. doi.org

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.

Industrial, Urban & Environmental Sustainability



Jonathan Platt
Mechanical Engineering, B.S.
Spring 2028
Vanderbilt University

Silicon solar cells waste much of sunlight's blue and ultraviolet light as heat, and their efficiency has nearly reached its theoretical ceiling. One way past that ceiling is a thin coating that turns the wasted blue light into near-infrared light, which silicon converts best. The best coating materials found so far contain toxic lead. We studied a lead-free alternative: the crystal $Cs_2AgBiCl_6$ doped with ytterbium, because ytterbium naturally emits light at that near-infrared wavelength. No one had previously pushed this material's conversion efficiency past 20%.

We made thin films of the crystal by evaporating its ingredients onto a glass substrate in a vacuum chamber, then heated them at 250–300 °C in air or nitrogen to activate the ytterbium. We also made related compounds to test how the recipe changes the light they absorb. We used X-ray diffraction to confirm each film formed the correct crystal structure, measured which wavelengths the films emit, and mapped which absorbed wavelengths produce that emission, revealing whether the crystal passes its energy to the ytterbium. Finally, we measured photoluminescence quantum yield: the fraction of absorbed light re-emitted.

Films with ytterbium emit near-infrared light at 990 nm; films without it do not. That emission appears only when the crystal itself absorbs the light, and only after heating: the crystal collects the energy and hands it to the ytterbium. The best films observed were 50% efficient but faded below 20% within a month in air. Sealing them with a thin protective layer raised the efficiency to 75%, the highest ever reported for this material, and kept it stable for over four months.

These films are a stable, lead-free coating that could make silicon solar cells more efficient.

NYU TANDON SCHOOL OF ENGINEERING

Physical Vapor Deposition and Optical Characterization of Quantum Cutting materials for increasing solar cell efficiencies

Jonathan Platt,[1]

Abstract

Silicon solar cells turn blue and ultraviolet light into heat instead of electricity. One fix is a coating that turns one high-energy photon into lower energy photons, quantum cutting, which silicon can absorb. The quantum cutting films discovered contain lead. Our films are lead-free double perovskite films doped with ytterbium. These films are able to convert 75% of the light they absorb back out at 990 nm, the best result anyone has reported for this material. Thus, these films show great promise increasing silicon solar cell efficiency.

Background on the Project

- Silicon solar cells are close to their theoretical efficiency limit.
- A coating/film can help increase efficiency by absorbing one high-energy photon and giving back light closer to the energy silicon uses best.
- Ytterbium is an apt dopant as it emits at 990 nm, where silicon is most efficient.

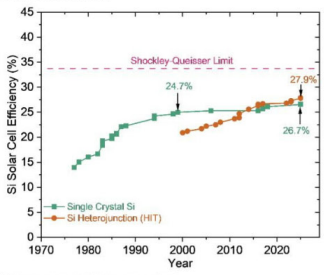


Figure 1. Silicon solar cell efficiency over time.

Materials and Methods

- Making the films: heat CsCl, AgCl, $BiCl_3$ and $YbCl_3$ until they evaporate and land on a glass slide inside a vacuum chamber.
- Annealing: 250 °C or 300 °C for one hour, either in nitrogen or in air.
- Testing: X-rays to check the crystal structure, and a fluorimeter to measure what light the films give off and how much.

Results I: Structure

- Replacing bromine with chlorine shrinks the crystal. The X-ray peaks shift steadily to the right, which means the swap happens evenly.
- Adding ytterbium dopant has little effect on the X-ray pattern, so it fits into the crystal without breaking it.

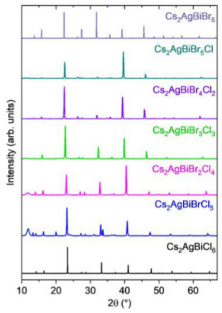


Figure 2. X-ray patterns as chlorine replaces bromine. The peaks shift right, showing the crystal shrinks smoothly.

Results II: Light Emission

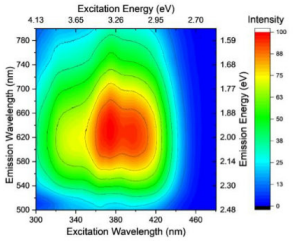


Figure 3. Emission-excitation map of an as-deposited Yb-doped film. Emission wavelength runs vertically, excitation wavelength horizontally, with color showing intensity. Emission peaks near 630 nm under ~375 nm excitation.

- Ytterbium doped films emit at 990 nm. Films without it do not, so the emission comes from the ytterbium.
- After annealing 53% of the absorbed light comes back out at 990 nm, but that drops to 15% after 40 days sitting in air.
- Adding a protective passivation layer raises it to 75% and holds it there for over four months.

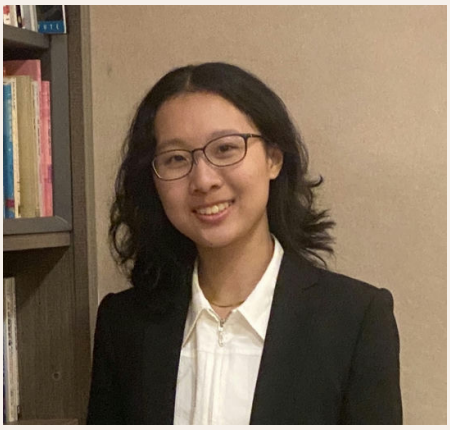
Works Cited

[1] P. Jain, E. Z. Tsang Xu, S. J. Paul, A. Morales, X. Liao, U. Rahman, S. Sarp and E. S. Aydil, paper in preparation.

Acknowledgements

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project. The author thanks Pulkita Jain for her mentorship and guidance throughout this project, and Professor Eray S. Aydil for the opportunity to work in his group.

Industrial, Urban & Environmental Sustainability



Kelsey Cheung
Mathematics, B.S.
Spring 2028
NYU Tandon

Molecular dynamics simulations provide a useful framework for studying the structure and mechanical behavior of amorphous materials. In this project, we investigate the inherent-state energy landscape of glass-forming system across a range of temperatures. Independent equilibrium configurations are generated using multiple simulation walkers and subsequently quenched through energy minimization to obtain their corresponding inherent states. For each minimized configuration, we measure the potential energy, virial pressure, deviatoric stress, stress magnitude, and the second stress invariant (J_2).

The resulting distributions are compared across temperature to determine how thermal sampling influences the regions of the energy landscape explored by the system. The simulations are being refined using standard Kob-Andersen binary-mixture parameters, longer equilibration runs, and additional temperatures near the apparent crossover region. This work establishes a computational pipeline for connecting finite-temperature molecular dynamics trajectories with the underlying inherent-state landscape and supports future studies of glassy relaxation and steady plastic flow.

NYU TANDON SCHOOL OF ENGINEERING

Kelsey Cheung
Advisor: Prof. Phillip Rauscher

Temperature Dependence of Inherent - State Energy and Pressure in a Kob-Anderson Glass

Background and Motivation

Amorphous and polymer glasses exhibit complex mechanical behavior because their properties depend strongly on temperature, preparation history, aging, and deformation. Molecular dynamics simulations provide atomistic information about these materials, but they are limited to relatively short length and time scales.

The broader project aims to compare the statistical and thermodynamic properties of equilibrium, quiescent, and steadily flowing glasses. In particular, distributions of inherent-state energy, stress, modulus, and dynamics may help characterize how glassy structure changes during aging, yielding, and plastic flow.

The present work establishes a temperature-dependent baseline for a quiescent model glass. These results can later be compared with configurations sampled during steady plastic flow.

Inherent States

An inherent state is the local energy minimum obtained by minimizing a sampled configuration.[1] Minimization removes thermal motion, revealing the underlying glass structure and any residual internal stress.

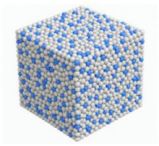


Figure 1. snapshot of the Kob–Andersen glass

Research Question

How does preparation temperature affect the inherent-state energy, pressure, and residual stress distributions of a model amorphous material?

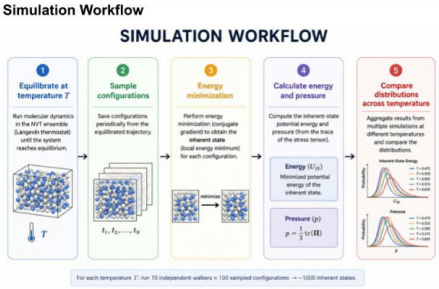
Model system

Lennard-Jones reduced units
Fixed simulation-cell volume
Molecular dynamics performed using LAMMPS
Multiple independent trajectories used for statistical sampling

Research Methods

1. Construct the Kob–Andersen system.
2. Equilibrate the system at a selected preparation temperature.
3. Run multiple independent trajectories.
4. Save configurations at regular time intervals.
5. Minimize each saved configuration at fixed volume.
6. Measure the potential energy and stress tensor of each minimized configuration.
7. Combine the results from all walkers.
8. Compare the resulting distributions across temperatures.

Simulation Workflow



Sampling (Each Temperature)

- 10 independent walkers
- 100 configurations sampled per walker
- 1,000 inherent states per temperature
- Configurations sampled at regular intervals along each trajectory

Temperature Set

T=0.4, 0.425, 0.45, 0.475, 0.5, 0.525, 0.55, 0.575, 0.6, 0.65, 0.7, 0.75, 0.8, 1

Measured Quantities

1. **Inherent-State Energy**

The minimized potential energy measures the depth of the local energy basin reached by each sampled configuration.

2. **Pressure**

$$p = \frac{1}{3}\mathrm{tr}(\mathbf{\Pi}) = \frac{\Pi_{xx} + \Pi_{yy} + \Pi_{zz}}{3}$$

Research Findings

- Inherent-state distributions vary systematically with temperature.
- Lower-temperature systems generally access deeper energy minima.
- Mean inherent-state energy increases monotonically with preparation temperature, indicating that colder systems sample deeper regions of the potential-energy landscape.
- Mean inherent pressure also increases monotonically with preparation temperature. Lower-temperature configurations minimize to states with more negative internal pressure, while the width of the pressure distribution remains approximately constant.
- Mean inherent pressure and mean inherent-state energy show a strong linear correlation across temperatures.

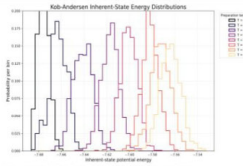


Figure 2. Combined inherent-state energy distributions for different temperatures.

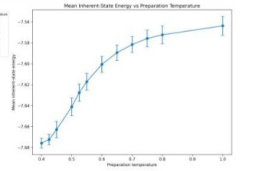


Figure 3. Mean inherent energy versus temperature

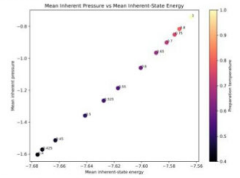


Figure 3. Mean inherent pressure versus mean inherent-state energy.

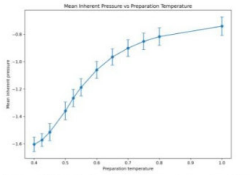


Figure 4. Mean inherent pressure versus temperature

Conclusions and Next Steps

Preparation temperature strongly influences the inherent structures sampled by the Kob–Andersen model glass. Colder systems minimize into deeper-energy states with more negative internal pressure, while warmer systems sample progressively shallower regions of the potential-energy landscape.

Future work will:

- Complete the analysis of residual-stress distributions.
- Compare the quiescent baseline with configurations sampled during steady plastic flow.
- Examine whether flow produces a reproducible nonequilibrium distribution of inherent-state properties.
- Investigate the relationship between inherent-state structure and mechanical response.

Works Cited

[1] Stillinger, Frank H., and Thomas A. Weber. "Hidden Structure in Liquids." Physical Review A, vol. 25, no. 2, 1982, pp. 978–989.

Acknowledgments

I thank the NYU Tandon School of Engineering Office of Undergraduate Academics for generously funding this project. I also thank Professor Phillip Rauscher for his guidance and support.

Industrial, Urban & Environmental Sustainability



Luyu Zhao
Computer Science, B.S.
Spring 2027
NYU Tandon



Yixia Yu
Computer Science, B.S.
Fall 2027
NYU Shanghai

Understanding GPU Power Dynamics in Large Language Model Inference

In recent years, the use of AI has exploded, leading to a rapidly increasing demand for data center GPUs, raising both environmental and logistical concerns due to the associated energy consumption. Especially from a data center operator and grid planning perspective, fine-grained power models have therefore become increasingly valuable, for example for grid incorporation, provisioning and capacity planning, all factors that need to be taken into account when trying to make data centers more energy efficient. Although model training is highly energy intensive, the large volume of daily inference requests means that AI inference may account for a substantial share of future AI computing demand. This study therefore investigates the power behavior of individual LLM inference requests and the system and workload factors that shape their power profiles. We developed a measurement framework that records continuous GPU power traces throughout the complete inference cycle, including idle, prefill, and decode phases. Experiments were conducted across a broad range of hardware platforms, language models, prompt categories, batch settings, parallelism strategies, and precision and quantization methods. Our results show that GPU power consumption varies considerably across models, hardware platforms, and serving configurations, while the effect of prompt content is generally much smaller. We also observed that initialization behavior, model optimization, and runtime settings can influence both power characteristics and overall inference efficiency. The collected measurements enable us to characterize workload-specific power distributions rather than relying only on aggregate averages. These distributions further support the development of stochastic models for estimating aggregate power demand in data centers running large-scale AI inference workloads, providing a foundation for more realistic energy evaluation and infrastructure planning.

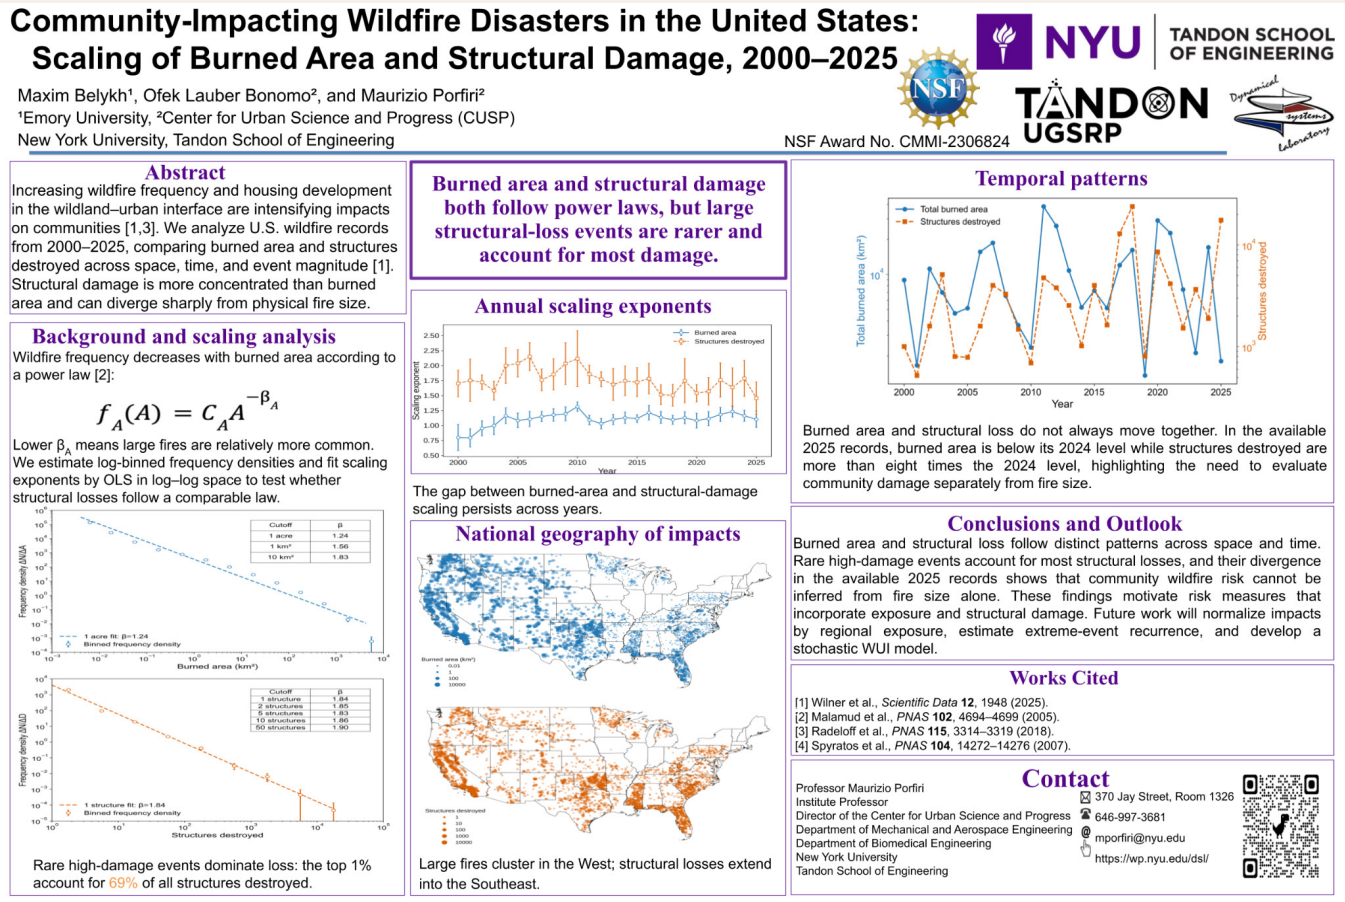
Industrial, Urban & Environmental Sustainability



Maxim Belykh
Economics, B.A.
Data Science, B.S.
Spring 2026
Emory University

Spatial and Temporal Scaling of Community-Impacting Wildfires in the United States

Wildfires that reach or threaten communities vary in both physical size and damage to the built environment, yet these dimensions are rarely analyzed together at the national scale. This study characterizes the spatial and temporal structure of community-impacting wildfires across the United States from 2000 to 2025, focusing on burned area and the number of structures destroyed. We use a harmonized geospatial wildfire dataset containing event locations, burn-zone area, structural losses, and ignition year. Wildfires are mapped, with symbol size representing both burned area and structures destroyed. We estimate frequency–magnitude relationships for burned area and structural damage using logarithmically binned distributions and power-law scaling models. To assess robustness, we repeat the scaling analysis across alternative lower cutoffs for event size and damage. Spatial analyses compare the geographic concentration of large and highly destructive fires, while temporal analyses examine annual variation nationally and separately by state. The results show that both burned area and structural losses are heavy-tailed. National frequency–area scaling is close to an inverse power law with an exponent near one, indicating that very large fires occur more often than under thin-tailed distributions. Structural damage is even more concentrated: most fires destroy few buildings, while a small fraction accounts for the majority of total losses. These relationships remain broadly consistent across cutoff choices, although estimated exponents vary with the lower threshold. Spatial and temporal analyses reveal substantial regional clustering, year-to-year variation, and meaningful differences among states. Together, these findings show that wildfire risk cannot be summarized by average event size alone. Integrating burned area, structural destruction, spatial concentration, and temporal variation provides a more complete framework for understanding community wildfire impacts and identifying regions most exposed to extreme events.



Industrial, Urban & Environmental Sustainability



Morgan Huzjak
Chemical and Biomolecular Engineering, B.S.
Spring 2027
Michigan State University

On the Role of Shunt Resistance on the Performance of Underwater Solar

Underwater solar power offers a way to run submerged devices without needing to surface, but its viability depends on how well solar cells perform under water's dimmed, narrowed light spectrum. Shunt resistance, a measure of current lost to leakage paths (shunts) caused from microscopic internal defects, is known to be detrimental to outdoor solar cell efficiency, specifically at low light intensities, but its role on underwater solar is less clear. In this project, we examined how shunt resistance influences efficiency at varying depths and water clarity. Using a detailed-balance model, three established photovoltaic materials, silicon, cadmium telluride, and amorphous silicon, were modeled across a range of shunt resistance values and depths of 0 to 10 meters, at three locations with different water clarity: Key West, Florida; Roatan Island, Honduras; and the South Pacific. Results showed that shunts' effect on efficiency grows slightly in murkier water, but even so, efficiency at literature-reported average shunt resistance values differed only marginally from efficiency at a theoretical infinite shunt resistance, across all three materials and locations. These findings suggest that for well-established solar cell materials, shunts are likely not a major limiting factor for underwater performance. However, recent findings suggest that newer, less-optimized photovoltaic materials are more sensitive to shunts, meaning that improving shunt resistance in these and other emerging materials could yield a larger performance benefit.

NYU

TANDON SCHOOL OF ENGINEERING

The Role of Shunt Resistance in Underwater Solar Materials

Morgan Huzjak, Jason A. Röhr

General Engineering, Tandon School of Engineering, New York University

Abstract

Underwater solar power offers a way to run submerged devices without needing to surface, but its viability depends on how well solar cells perform under water's dimmed, narrowed light spectrum. Shunts, a measure of current lost to leakage paths caused from microscopic internal defects, is known to be detrimental to outdoor solar cell efficiency, but its role on underwater solar is less clear. We examined how shunt resistance influences efficiency in varied underwater environments. various photovoltaic materials were modeled across a range of shunt resistance values. We show that shunts' effect on efficiency grows slightly in murkier water, but even so, efficiency at literature-reported average shunt resistance values differed only marginally from efficiency at a theoretical infinite shunt resistance. These findings suggest that for well-established solar cell materials, shunts are likely not a major limiting factor for underwater performance. However, recent findings suggest that newer, less-optimized photovoltaic materials are more sensitive to shunts, meaning that improving shunt resistance in these and other emerging materials could yield a larger performance benefit.

Background

- Silicon dominates terrestrial photovoltaics, but its narrow ~1.1 eV bandgap makes it poorly suited underwater.
- Water strongly absorbs infrared & red light while transmitting blue-green light, so much of the spectrum Si is designed to harvest never reaches the cell at depth [3].
- As depth increases, the solar spectrum underwater narrows, shifting the optimum bandgap for maximum efficiency to ~1.8 eV - 2.2 eV.



Figure 1. Underwater Solar Powered Vehicle [5]

- Real-world solar cell performance is limited by non-ideal device behavior, particularly shunt resistance, a detrimental leakage pathway, caused by defects or impurities in the device, that allows current to bypass the external circuit [2]. Shunt's impact grows as light intensity decreases. Because underwater irradiance is inherently low-intensity and narrowed, low shunt resistance could disproportionately degrade underwater cell performance in a way not captured by standard analyses [1],[2].
- This project builds on the framework for underwater solar cells by explicitly incorporating shunt resistance as a variable to determine whether it meaningfully limits underwater photovoltaic performance.

Methodology

- These depth-resolved irradiance spectra served as the P_in inputs to the efficiency calculations in the model.

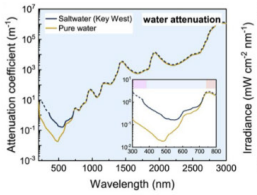


Figure 2. Attenuation coefficient spectra of salt water and pure water [4]

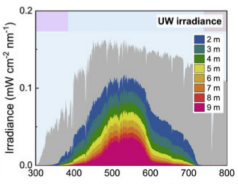


Figure 3. Underwater irradiance spectra [4]

Efficiency was calculated using:

$$\eta = P_{out} / P_{in}$$

$$P = I \times V$$

P_out was taken as the maximum power point (peak of this curve).

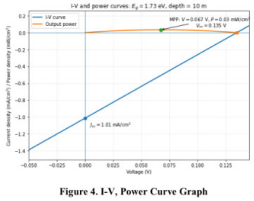


Figure 4. I-V, Power Curve Graph

Results

-At the surface, both established materials (Si, CdTe, a-Si) and emerging materials (MAPbI3, P3HT:PCBM) track close to the theoretical limit.

-At 10m depth, this gap widens substantially: established materials remain near the limit, but MAPbI3 and especially P3HT:PCBM show much larger efficiency losses, with performance dropping sharply as shunt resistance decreases.

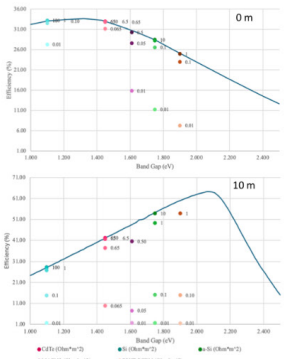


Figure 4. Efficiency as a Result of Bandgap at 0 m depth (top) and at 10 m depth (bottom)

Future Work

- Extend to newer, less-established photovoltaic materials.
- Incorporate external radiative efficiency (ERE) into the model to allow non-radiative recombination losses to be considered.

Works Cited

[1] J. A. Röhr, J. Lipton, J. Kong, S. A. Maclean, and A. D. Taylor, *Joule*, vol. 4, pp. 840–849, 2020.

[2] C. M. Proctor and T.-Q. Nguyen, *Applied Physics Letters*, vol. 106, no. 8, 083301, 2015.

[3] J. A. Röhr, B. E. Sartor, J. Lipton, and A. D. Taylor, *Nature Photonics*, vol. 17, pp. 747–754, 2023.

[4] J. A. Röhr, B. E. Sartor, J. N. Duenow, et al., *iScience*, vol. 25, 104531, 2022.

[5] https://eepower.com/tech-insights/next-gen-solar-power-for-underwater-vehicles/

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.

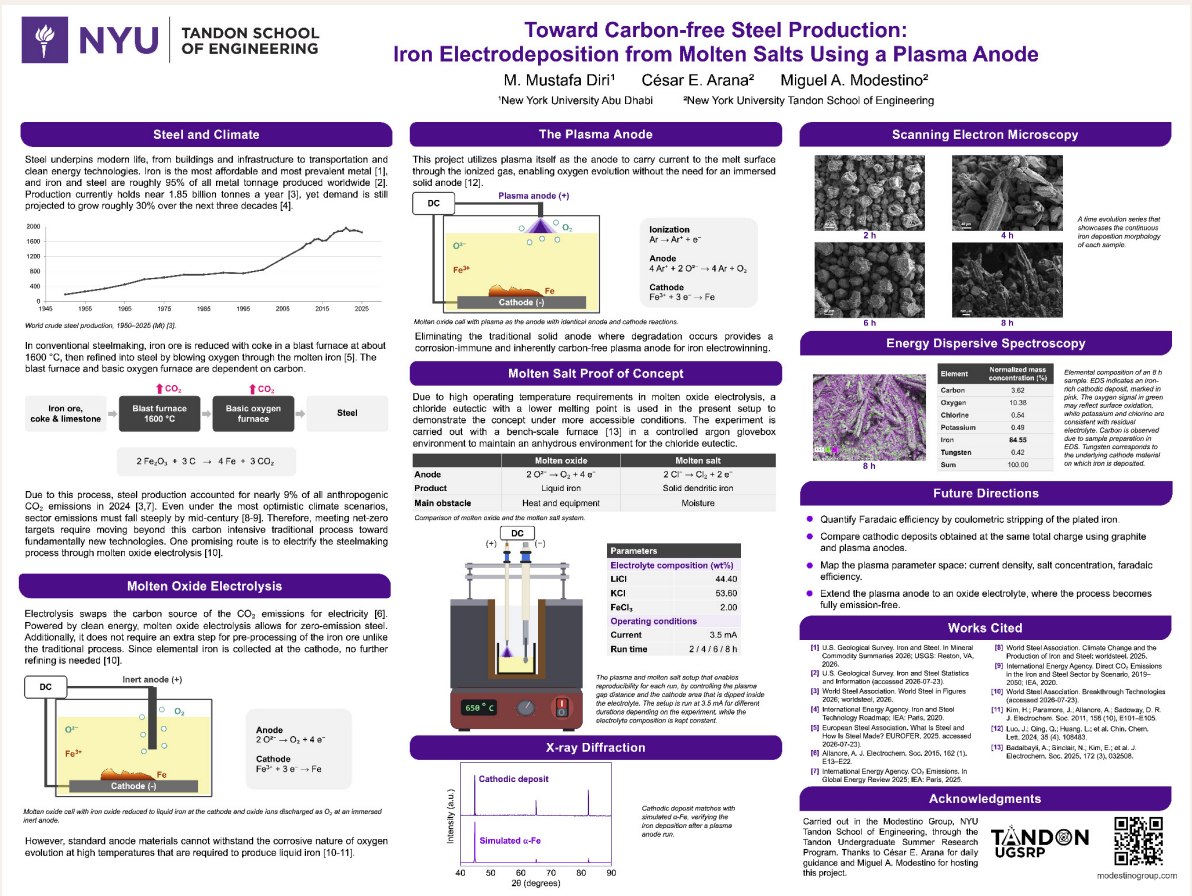
Industrial, Urban & Environmental Sustainability



Muhammet Mustafa Diri
Mechanical Engineering, B.S.
Spring 2027
NYU Abu Dhabi

Toward Carbon-free Steel Production: Iron Electrodeposition from Molten Salts Using a Plasma Anode

Steel production has surged within the last few decades to underpin modern life, from buildings and infrastructure to transportation and clean energy technologies. The demand for steel is expected to grow roughly 30% over the next three decades. However, in 2024, steel production accounted for nearly 9% of all anthropogenic CO2 emissions. Even under the most optimistic climate scenarios, meeting the net-zero targets therefore requires moving beyond the conventional blast furnace and basic oxygen furnaces, which are stoichiometrically dependent on carbon, toward fundamentally new technologies. One promising route is to electrify the steelmaking process through molten oxide electrolysis (MOE), which swaps the carbon source of the CO2 emissions for electricity. Powered by clean energy, MOE allows for fully sustainable, zero-emission steel. However, standard anode materials cannot withstand the corrosive nature of oxygen evolution at high temperatures that are required to produce liquid iron. Instead, this project utilizes plasma itself as the anode to carry current to the melt surface through the ionized gas, enabling oxygen evolution without the need for an immersed anode. By eliminating the traditional solid anode where degradation occurs, plasma provides a corrosion-immune and inherently carbon-free anode for iron electrowinning. As a proof of concept, iron is deposited from a chloride eutectic that allows the experiment to run in a lower-temperature environment with a bench-scale furnace in a controlled argon glovebox environment. Deposits are characterized by XRD to identify crystalline phases, while their evolution with deposition time is assessed by SEM and EDS to evaluate morphology and elemental composition. Faradaic efficiency is quantified by coulometric stripping of the iron plate through a potentiostat. Demonstrating iron deposition under a plasma anode in molten salt systems validates the novel electrode concept for extension to oxide electrolytes, where the process becomes fully emission-free.



Industrial, Urban & Environmental Sustainability



Nicholas Reinoso
Electrical Engineering and Computer Engineering (dual major), B.S.
Spring 2029
NYU Tandon

FloodNet: Sustaining Energy Neutral Operation Through Server-Side Transmission Control

FloodNet operates 450 low-cost, solar-powered ultrasonic sensors that report street-level flooding across New York City's five boroughs. Each sensor measures water level every minute and transmits over cellular at a reporting interval that can be changed remotely, from every minute (the default) up to every 30 minutes. In winter, shorter days and building shading cut solar harvest at shaded sites below what a sensor consumes. Batteries drain over weeks and sensors brown out until field crews visit to swap them. To determine how much power slowing the reporting rate actually saves, bench measurements were taken on hardware. A FloodNet sensor was instrumented with a Qoitech Otii Arc power analyzer at the battery's 3.60 V nominal voltage, recording current at 4,000 samples per second across 1-, 15-, and 30-minute reporting schedules in 63-, 23-, and 16-hour captures. The analyzer's UART logging tied each current spike to a logged firmware event. Slowing reporting from 1 to 15 minutes cut average current draw by 41 percent, from 2.82 mA to 1.66 mA, stretching a full charge of the sensor's Samsung INR18650-35E lithium-ion cell from 49 to 84 days with no solar input. Slowing further to 30 minutes saved nothing measurable, because the firmware's 900-second MQTT keepalive wakes the cellular modem halfway through each 30-minute interval to keep its server connection alive: the radio wakes four times per hour on either schedule. Fifteen-minute reporting therefore delivers the most live data at the sensor's lowest measured power. A server-side reporting policy is proposed, not yet deployed: sensors above a battery voltage threshold keep 1-minute reporting, sensors below it drop to 15 minutes, and a forecast storm or nearby flooding restores 1-minute reporting so flood coverage is never compromised.

NYU TANDON SCHOOL OF ENGINEERING

NYU Center for Urban Science + Progress

Powering FloodNet: Optimizing Sensor Reporting for Battery Life

15-Minute Reporting Cuts Sensor Energy Use 41% and Extends Battery Life from 49 to 84 Days

Nicholas Reinoso,[1] Charlie Mydlarz,[2] Andrea I. Silverman[2,3]

1. Undergraduate Summer Research Program, NYU Tandon School of Engineering 2. Center for Urban Science and Progress, NYU Tandon 3. Department of Civil, Urban, and Environmental Engineering, NYU Tandon

Project Supervisors: Charlie Mydlarz and Andrea I. Silverman

Acknowledgments: The author thanks the NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of this project. Thanks also to Praneeth Challagonda, Chris Hoyte, Amanpreet Kaur, Polly Pierone, and Bea Steers.

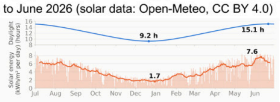
Abstract

FloodNet is a street-level flood monitoring network with 450 sensors across New York City's five boroughs. Each solar-powered unit measures water level every minute by ultrasonic sensor and reports over cellular at a remotely settable reporting interval (default 1 minute). In winter, short days and building shading cut solar harvest below what a shaded sensor consumes. Batteries run down over weeks, sensors shut off and stop reporting, and each one stays dark until a crew spends staff time on a site visit, swapping the sensor for a charged one. To measure how much energy each reporting interval actually consumes, a FloodNet sensor's current draw was recorded with a Qoitech Otii Arc power analyzer at 1-, 15-, and 30-minute intervals (setup below). Slowing reporting from 1 to 15 minutes cut average draw by 41%, from 2.82 to 1.66 mA, stretching a full charge of its Samsung INR18650-35E cell (3,350 mAh) from 49 to 84 days with no solar input (Figure 2). Slowing to 30 minutes saved nothing more: a 900-second MQTT keepalive wakes the radio mid-interval anyway. Fifteen-minute reporting delivers the most live data at the lowest measured power, at the price of reports up to 15 minutes late. A server-side policy is proposed: 15-minute reporting below a battery voltage threshold, with 1-minute reporting restored when storms or nearby flooding threaten.

Testing

Testing used a Qoitech Otii Arc power analyzer recording current at 4,000 samples per second at the battery's 3.60 V nominal voltage. The 1-, 15-, and 30-minute reporting intervals were captured for 63, 23, and 16 hours, respectively. The Arc's serial (UART) log matched each current spike to a logged sensor event.

Figure 1: Daylight and daily solar energy (kWh/m² per day) near NYC, July 2025 to June 2026 (solar data: Open-Meteo, CC BY 4.0)



Results

Figure 2: Predicted state of charge at 1-, 15-, and 30-minute reporting on the Samsung INR18650-35E cell, no solar input

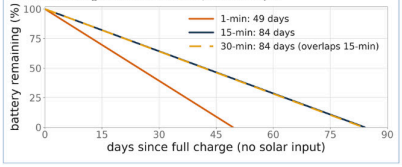


Figure 3: Radio duty cycle at each reporting interval: the share of each hour the radio is on (dark pulses are uploads, orange are keepalive pings)

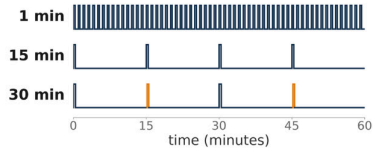
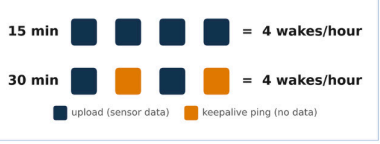


Figure 3 shows the share of each hour the radio is on. At the default 1-minute interval the radio consumes the majority of the battery, and radio-on time is the part of the energy budget the reporting interval controls. That on-time is nearly identical at 15 and 30 minutes, which is why slowing past 15 minutes saves nothing measurable.

The averages behind Figures 2 and 3 come from deliberately long captures: 63, 23, and 16 hours at the three reporting intervals. Averaging over dozens of reporting cycles and many hours of cellular network variability makes the averages stable rather than dependent on any single stretch of conditions, computed over complete hours only with the first boot hour excluded.

The radio wakes only once per reporting interval to send its stored readings, but the MQTT connection times out after 900 seconds without traffic, so the 30-minute reporting interval adds a mid-interval keepalive ping (orange pulses, Figure 3).

Figure 4: Radio wakes per hour at 15- and 30-minute reporting



Because of the keepalive ping (Figure 4), 15- and 30-minute reporting produce the same four radio wakes per hour and nearly identical radio-active time, so their energy use is equivalent (Figure 3). Radio wakes draw the sensor's highest instantaneous current (transmit peaks above 100 mA against a 1.26 mA idle floor). The added keepalive wake cancels the saving that halving the upload rate should provide.

The cost of slower reporting is latency: a report of street flooding can arrive up to the reporting interval late. Proposal (1) below counters this by restoring 1-minute reporting when a storm is forecast or nearby flooding is detected.

Future Work

Proposals: (1) a server-side policy: 1-minute reporting above a battery voltage threshold, 15-minute below, with 1-minute reporting restored when storms or nearby flooding threaten; (2) extending the MQTT keepalive from 15 minutes (900 s) to 19 minutes (1,140 s), just under the Azure IoT Hub limit, and measuring energy at the longer setting; (3) shortening the 250 ms pause the firmware leaves between queued records during an upload.

Industrial, Urban & Environmental Sustainability



Nour Al Saqqa
Civil Engineering, B.S.
Spring 2027
NYU Abu Dhabi

Planning for Quality-of-Life Environmental Monitoring at the Manhattan Detention Complex Construction Site for Residents of Chinatown

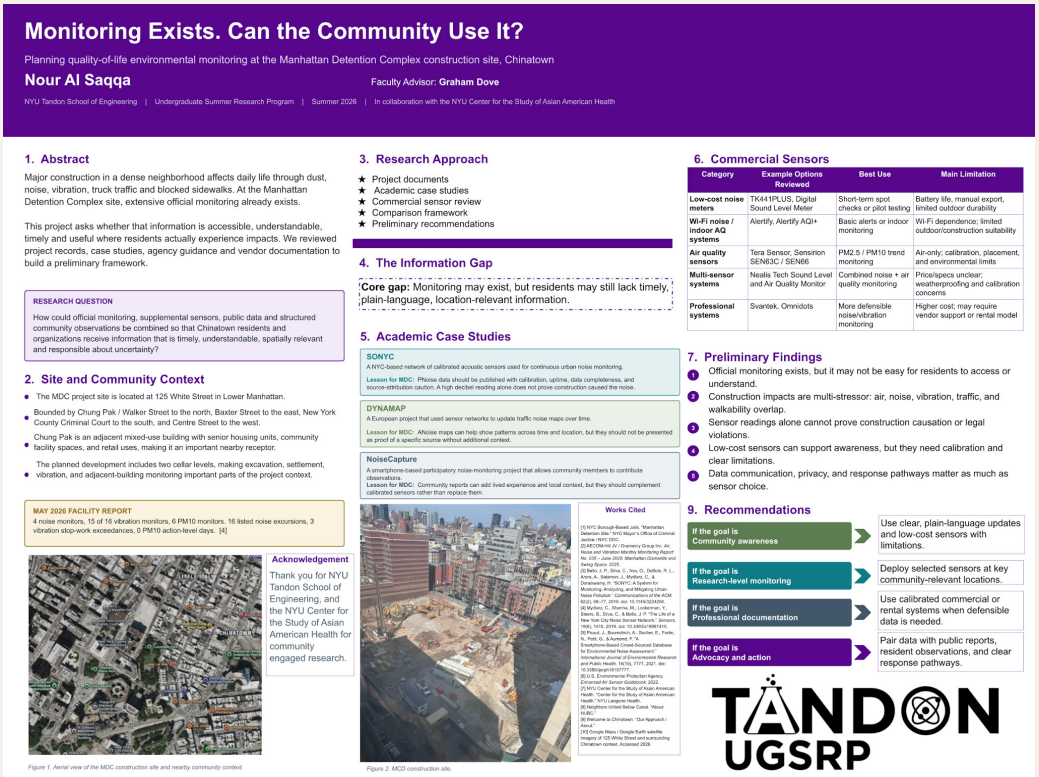
Major construction projects in dense urban neighborhoods can create environmental, public health, and quality-of-life concerns for nearby communities, including dust, diesel emissions, noise, vibration, truck activity, sidewalk disruption, and reduced access to local businesses and community spaces. This project examines the Manhattan Detention Complex construction site at 125 White Street in Chinatown, Lower Manhattan, and asks how environmental monitoring could better support residents, small businesses, and community organizations responding to long-term construction activity.

Conducted in collaboration with the NYU Center for the Study of Asian American Health, this research is informed by community-engaged approaches to health equity and by local advocacy and civic engagement efforts, including the work of Neighbors United Below Canal and Welcome to Chinatown. The objective is to investigate how existing construction-related monitoring, low-cost sensors, commercial environmental monitoring systems, community observations, and public data could be combined to provide clearer and more accessible information about construction-related impacts.

The research uses a literature review and framework-development approach. I reviewed site context, community concerns, academic case studies in urban environmental sensing, and commercially available monitoring options for air quality, noise, vibration, and multi-sensor construction monitoring.

Preliminary findings suggest that official monitoring may address technical concerns such as groundwater, dewatering, structural movement, vibration, and stormwater compliance. However, a gap may remain between compliance-oriented monitoring and the information residents need to understand daily conditions near homes, sidewalks, businesses, senior-serving facilities, and other community locations. Low-cost and commercial sensors may help make impacts more visible, but their usefulness depends on calibration, placement, maintenance, data access, privacy protections, and careful communication of uncertainty.

This project contributes to a preliminary framework for evaluating supplemental monitoring options that can support community awareness, advocacy, and responsible environmental data communication.



Industrial, Urban & Environmental Sustainability



Shely Dash
Computer Science, B.S.
Fall 2029
Purdue University

Interpretable Machine Learning for Planning Fixed and Mobile Electric Vehicle Charging Infrastructure

The rapid adoption of electric vehicles (EVs) has created new challenges for transportation planners seeking to deploy charging infrastructure that balances user accessibility, operational efficiency, and economic sustainability. Although transportation simulations can evaluate different charging network designs, they are computationally expensive and limit the number of scenarios that can be explored. This research investigates whether interpretable machine learning can provide fast, reliable predictions while also explaining how key planning decisions influence charging network performance.

Using a simulation of the Anaheim transportation network, I analyzed hundreds of charging infrastructure scenarios that varied in the number, type, and capacity of charging facilities. For each scenario, I developed features describing local travel demand, road network characteristics, nearby competition, and charging infrastructure. These features were used to train Neural Additive Models (NAMs) to predict revenue and average waiting time for both fixed and mobile charging facilities. To evaluate performance, the NAM was benchmarked against Ridge Regression, Random Forest, XGBoost, and Multilayer Perceptron models using RMSE, MAE, and R^2.

Preliminary results indicate that ensemble learning methods achieve the highest predictive accuracy, while the Neural Additive Model provides competitive performance with the added advantage of interpretability. Ongoing work focuses on refining the feature set and generating feature-response curves that quantify how changes in demand, competition, accessibility, and infrastructure characteristics affect revenue and user waiting times.

Beyond improving prediction speed, this work demonstrates how interpretable machine learning can support engineering design and economic decision-making for EV charging infrastructure. By revealing the marginal impact of key planning variables on operational performance and financial outcomes, the framework provides transportation agencies and charging network operators with transparent tools for evaluating infrastructure investments, allocating resources efficiently, and designing charging systems that balance profitability with service quality as EV adoption continues to grow.

NYU TANDON SCHOOL OF ENGINEERING | C2SMART CONNECTED CITIES WITH SMART TRANSPORTATION

Interpretable Machine Learning for Joint Development of Fixed and Mobile EV Charging Facilities

Shely Dash[1], Wenjia Shen[2], Joseph Chow, Ph.D., P.E.[2]
[1]Barnard College, Columbia University, [2]Department of Civil and Urban Engineering, NYU Tandon School of Engineering

BACKGROUND

Electric vehicle (EV) adoption has increased the demand for efficient charging infrastructure. Transportation simulation accurately evaluates charging strategies but is computationally expensive. In this study, we used Neural Additive Models (NAMs) and ensemble learning models to predict charging network revenue and waiting time across 354 simulated scenarios while maintaining model interpretability.

MATERIALS AND METHODS

The Anaheim transportation network was used to generate 354 EV charging scenarios with varying facility locations, capacities, and demand patterns. Five machine learning models (NAM, Ridge Regression, Random Forest, XGBoost, and MLP) were trained using a reproducible 70/30 train-test split and evaluated with RMSE, MAE, and R^2.

SIMULATION & MODELING WORKFLOW

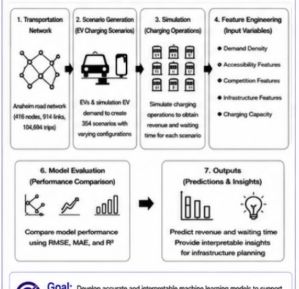


Figure adapted from TransportationNetworks (Anaheim Network). Available: https://github.com/bstabler/TransportationNetworks

NEURAL ADDITIVE MODEL (NAM)

$$\hat{y} = \beta_0 + \sum_{j=1}^{K} f_j(x_j)$$

NAM assumes the prediction is a sum of smooth functions, where each feature contributes independently.

EXAMPLE FOR REVENUE PREDICTION

$$\text{Revenue} = f(\text{Demand Density}) + f(\text{Accessibility}) + f(\text{Competition}) + f(\text{Infrastructure}) + \beta_0$$

Each $f_j(x_j)$ shows how that feature influences the outcome, allowing interpretable, feature-level insights.

TABLE 1. PREDICTIVE PERFORMANCE COMPARISON

Prediction Task (Predicting)	MODEL				
	Ridge	Random Forest	XGBoost	MLP	NAM (Neural Additive Model)
Fixed Revenue	High	High	**Best**	Moderate	High
Fixed Wait	High	High	**Best**	Moderate	High
Mobile Revenue	High	High	**Best**	Moderate	High
Mobile Wait	High	High	**Best**	Moderate	High

TABLE 2. MODEL CHARACTERISTICS

Characteristic	MODEL				
	Ridge	Random Forest	XGBoost	MLP	NAM (Neural Additive Model)
Accuracy (overall)	High	High	**Best**	Moderate	High
Interpretability	Moderate	Low	Low	Low	**Best**
Feature-level explanation	✓	×	×	×	✓

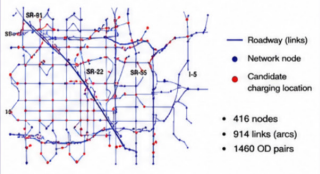
Table 1. Comparison of models' ability to predict charging network revenue and waiting time accurately. XGBoost achieved the highest predictive accuracy across all tasks, while NAM remained highly competitive and uniquely provided interpretable, feature-level explanations.

Legend: Best = highest accuracy; High = strong performance; Moderate = moderate performance; Low = limited interpretability/performance; ✓ = available; × = not available

DISCUSSION AND CONCLUSION

Machine learning substantially reduced the computational cost of evaluating EV charging infrastructure while maintaining strong predictive performance. Ensemble learning models achieved the highest accuracy, whereas NAM improved transparency by explaining how individual planning variables influenced revenue and waiting time. This framework provides transportation planners with a practical tool for rapidly evaluating charging infrastructure and supporting future optimization.

FIGURE 1. ANAHEIM TRANSPORTATION NETWORK



The Anaheim road network used to generate 354 charging infrastructure scenarios. Red dots indicate candidate locations where charging facilities can be deployed.

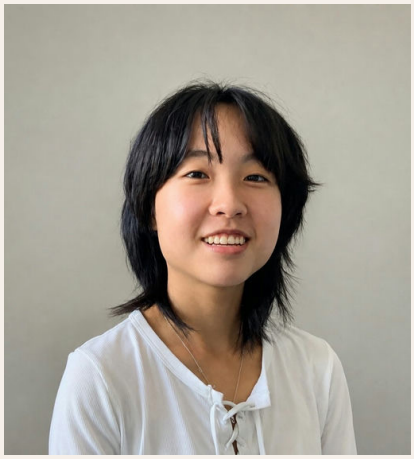
ACKNOWLEDGEMENTS

This research was conducted through the NYU Tandon Undergraduate Summer Research Program (UGSRP) under the mentorship of Wenjia Shen at the C2SMARTER Center. We gratefully acknowledge Wenjia Shen and Professor Chow for their guidance and the UGSRP program for their support.

WORKS CITED

[1] R. Agarwal et al., "Neural Additive Models: Interpretable Machine Learning with Neural Nets," *Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 21503–21518, 2021.
[2] L. Breiman, "Random Forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
[3] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, pp. 785–794, 2016.
[4] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
[5] V. H. L. Shen et al., "EVCS: A High-Fidelity Simulation Framework for Three-Sided Electric Mobility Markets," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 6, pp. 3316–3329, 2021.

Industrial, Urban & Environmental Sustainability



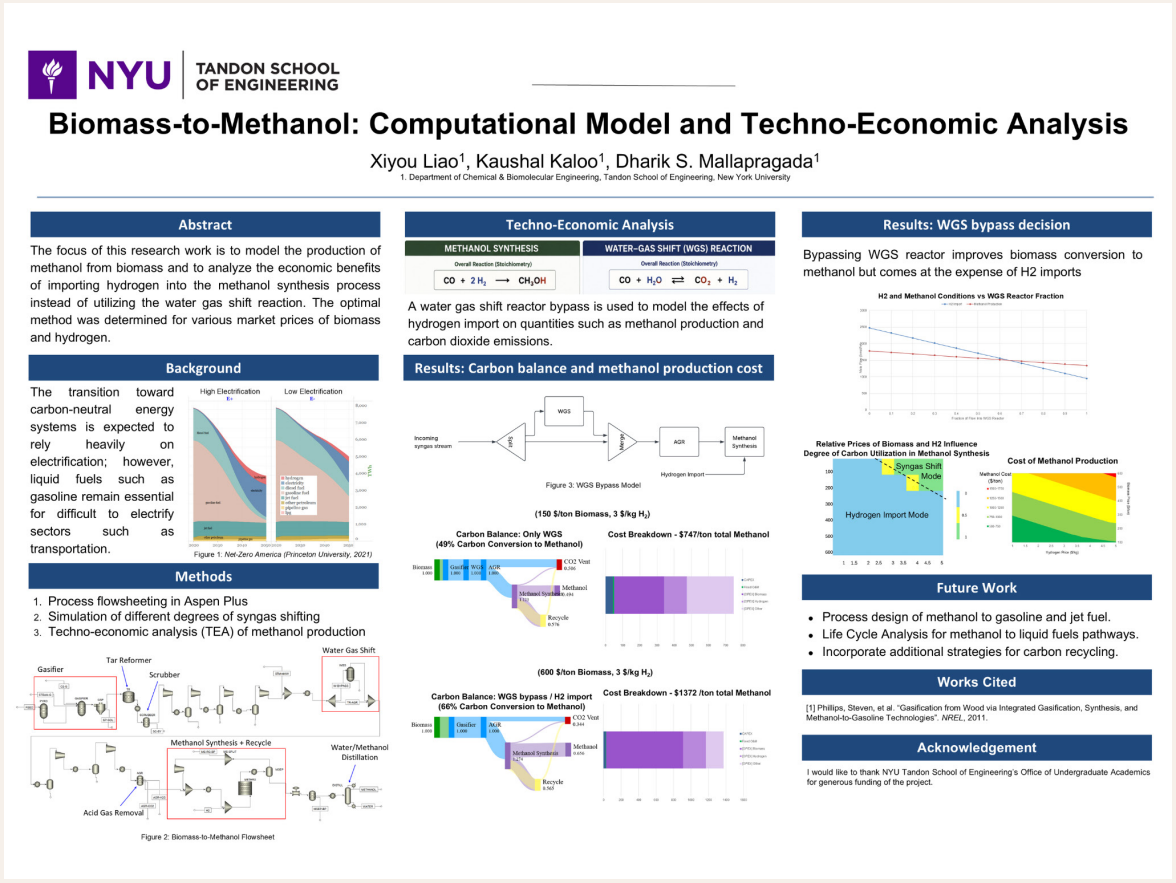
Ren Liao
Chemical and Biomolecular Engineering, B.S.
Spring 2027
NYU Tandon

Computational Modeling and Techno-Economic Analysis of Biomass to Methanol Process

The transition toward carbon-neutral energy systems is expected to rely heavily on electrification; however, liquid fuels such as gasoline remain essential for difficult to electrify sectors. Thus, developing sustainable pathways for liquid-fuel production is important for energy decarbonization. One method to achieve that goal is biomass-derived methanol because methanol serves as an intermediate for the production of gasoline and other hydrocarbon fuels.

This study models the biomass-to-methanol production process and evaluates the economic tradeoff between generating hydrogen through the WGS reaction and importing external hydrogen for methanol synthesis. An Aspen Plus V14 model was developed based on the NREL report Gasoline from Wood via Integrated Gasification, Synthesis, and Methanol-to-Gasoline Technologies. The process includes biomass gasification, syngas cleaning, water-gas shift (WGS), acid-gas removal (AGR), and methanol synthesis. The model was designed with the condition of achieving an H2/CO ratio of 2 before entering the methanol synthesis reactor.

A bypass upstream of the WGS reactor was varied between full WGS usage and complete WGS bypass with external H2 importation. A techno-economic analysis (TEA) was performed to calculate the favored WGS bypass fraction with varying biomass and hydrogen market prices. With increasing external hydrogen importation, more methanol is produced with a carbon-to-methanol conversion increase from 49% with full WGS utilization to 66% with only H2 supplementation. Complete hydrogen import is favored when hydrogen price is below $2.5/kg and biomass price is above $300/ton. Therefore, although the additional hydrogen importation increases feed costs, the resulting reduction in WGS-related capital and operating expenditures lowers the overall plant costs, demonstrating a potential economic advantage for external hydrogen integration in biomass-to-methanol production.





Lauren Wang
Data Science,B.S.
Spring 2027
NYU Shanghai

How Well Can Choice Models Recover Individual Preferences? A Ground-Truth Benchmark on EV Charging Choice

Discrete choice models predict an individual's choice among mutually exclusive alternatives based on each alternative's utility, and are central to modeling travel behavior such as EV charging station choice. They differ in how they represent individual-level taste heterogeneity and at what cost, but are usually compared by relative fit, which cannot indicate how close a model comes to the best achievable recovery when tastes are partly driven by unobserved factors. Toward informing model choice for large-scale EV charging data, we construct a synthetic benchmark with known individual parameters and a computable lower bound on recovery error, and evaluate four models in absolute terms. Synthetic charging station choice data follow Han et al. (2022): individual tastes are polynomial functions of income, time of day, and duration, with the cost coefficient fixed at −1 to place tastes in willingness-to-pay space. Two attribute-correlation conditions cross three sample sizes (N = 500/2,000/10,000). Adding 20% per-individual coefficient noise makes tastes non-deterministic in the covariates and yields an analytically computable irreducible error floor. MNL, MXL, AMXL, and TasteNet-MNL are compared on predictive fit, parameter recovery, value-of-time distributions, and computational cost. With covariates correctly specified and one observation per agent, TasteNet recovered travel-time coefficients to within 1.27× the floor at N = 10,000, against 2.16× for AMXL, and its error fell 22% as the sample grew while AMXL's stayed flat. On the weakly identified fast-charging coefficient, AMXL was worse than the population-mean baseline in all six conditions, and it took 12.4× longer than TasteNet at N = 10,000. For EV charging models, three implications follow: individual-level estimation offers little benefit for weakly identified attributes such as charging speed and is better left to fixed coefficients; agent-based methods need repeated observations per person, which stated-preference surveys lack but real GPS charging records provide; and their computation time grows faster than sample size, which matters at the scale of real data.

NYU · C2SMART CONNECTED CITIES WITH SMART TRANSPORTATION · BUILT @ NYU Behavioral Urban Informatics, Logistics, & Transport Lab · TANDON UGSRP

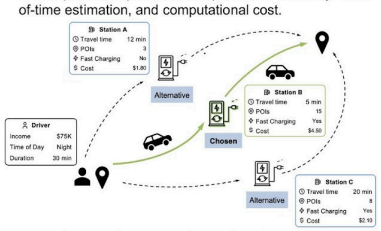
How Well Can Four Choice Models Recover Individual Taste Heterogeneity?

Fit, recovery, and cost on synthetic EV charging data

Lauren Wang[1], Xiyuan Ren[2], Joseph Chow[2]
[1]NYU Shanghai, [2]C2SMARTER, NYU Tandon School of Engineering

INTRODUCTION

Electric vehicle charging behavior varies substantially across individuals, making individual-level taste heterogeneity central to modeling charging demand. Existing discrete choice models represent heterogeneity in fundamentally different ways—from shared coefficients (MNL) and parametric distributions (MXL) to agent-based (AMXL) and neural (TasteNet-MNL) approaches—but their ability to recover true individual preferences remains unclear. Using synthetic EV charging data with known ground-truth parameters, we compare these four models across predictive performance, parameter recovery, value-of-time estimation, and computational cost.



Notes: ***POI****: number of nearby points of interest;* ***Cost****: charging price,* ***TT****: Travel Time,* ***DCFC****: DC Fast Charging. Driver covariates: income, time of day, charging duration.*

METHODS

Data design. Each session samples a driver with standardized covariates z_i (income, time of day, duration). Individual tastes are generated as:

$$\tilde{\beta}_{k,i} = a_{0,k} + a_{1,k}^{\mathrm{T}} z_i + z_i^{\mathrm{T}} A_{2,k} z_i + \varepsilon_{k,i}, \quad \varepsilon_{k,i} \sim \mathcal{N}(0, \sigma_k^2)$$

$$\beta_k(z_i)$$

where $k \in \{\mathrm{TT}, \mathrm{POI}, \mathrm{DCFC}\}$, with 20% Gaussian noise. Choices are generated by:

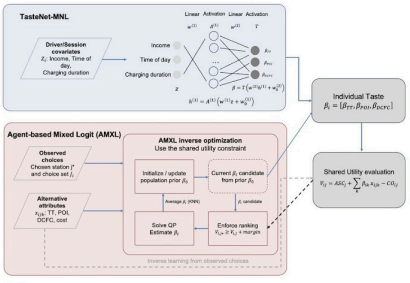
$$V_{ij} = \mathrm{ASC}_j + \sum_k \tilde{\beta}_{k,i}\, x_{ijk} - CO_{ij}, P(y_i = j) = \mathrm{softmax}(V_i)$$

Scenarios: $\{uncorrelated / TT\text{–}POI\ \rho = 0.6\} \times \{500, 2k, 10k\}$. Split = 5:1:1.

Models

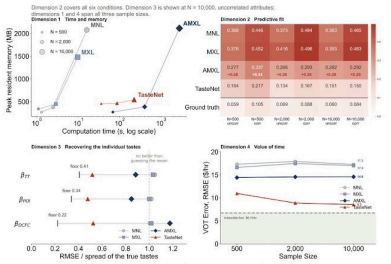
Model	How it represents preferences
MNL Multinomial Logit	The standard choice model: **weighted attributes**, highest utility wins. One set of weights for everyone, so differences between drivers fall into the error term.
MXL Mixed Logit	Weights vary randomly across drivers, drawn from a **distribution** the modeler assumes in advance, usually normal.
AMXL Agent-Based Mixed Logit	Works backwards, inferring the weights that would have made each observed choice **optimal**. One free set per driver, anchored to the population average, no distribution assumed.
TasteNet-MNL	A **neural network** maps driver characteristics to weights, so preferences vary smoothly with income, time of day, and duration.

Notes: The two models that estimate individual-level preferences.



RESULT

Four models across four dimensions



- **AMXL** achieves near-perfect training fit but generalizes poorly, indicating severe overfitting in the one-observation-per-driver setting.
- **TasteNet** recovers individual tastes most accurately, though a gap to the theoretical noise floor remains.
- **MNL** and **MXL** provide competitive prediction but are less effective at recovering individual-level preference heterogeneity.

DISCUSSION

Main finding. TasteNet-MNL provides the best overall balance between predictive accuracy, individual taste recovery, and computational efficiency. By learning preferences directly from observed covariates, it avoids restrictive distributional assumptions while generalizing substantially better than AMXL in sparse-data settings.

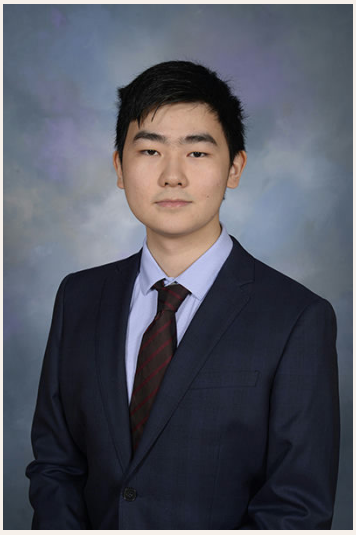
Limitations

- The synthetic data were generated from a smooth polynomial mapping, which naturally favors flexible function approximators.
- AMXL is evaluated in a one-observation-per-driver setting, outside its intended repeated-choice scenario.

References

[1] Ren, X., & Chow, J. Y. J. (2022). A random-utility-consistent machine learning method to estimate agents' joint activity scheduling choice from a ubiquitous data set. *Transportation Research Part B*, 166, 396–418.
[2] Han, Y., Pereira, F. C., Ben-Akiva, M., & Zegras, C. (2022). A neural-embedded discrete choice model: Learning taste representation with strengthened interpretability. *Transportation Research Part B*, 163, 166–186.

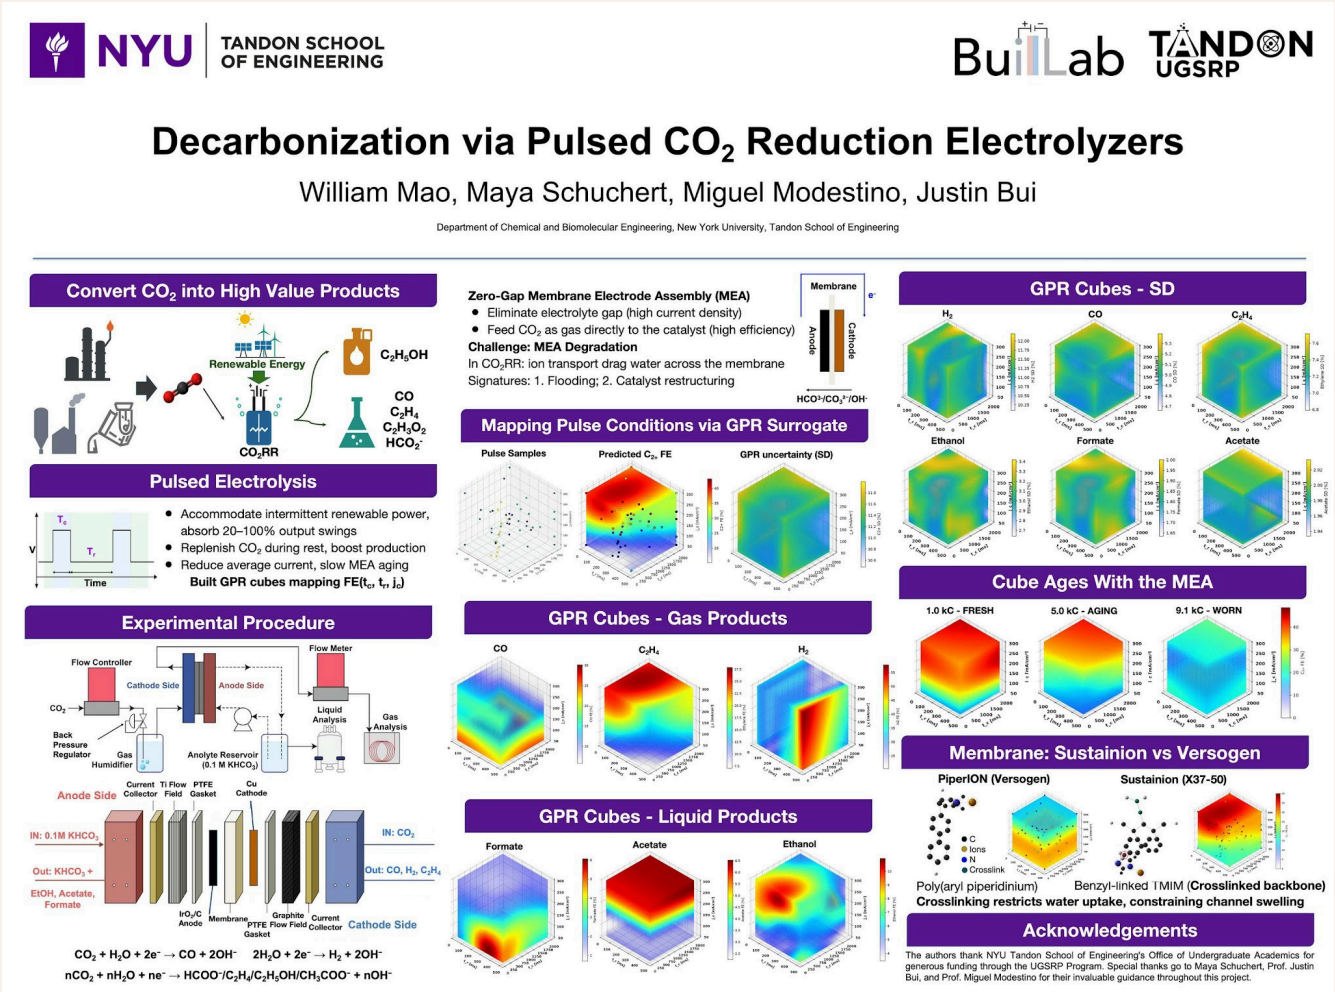
Industrial, Urban & Environmental Sustainability

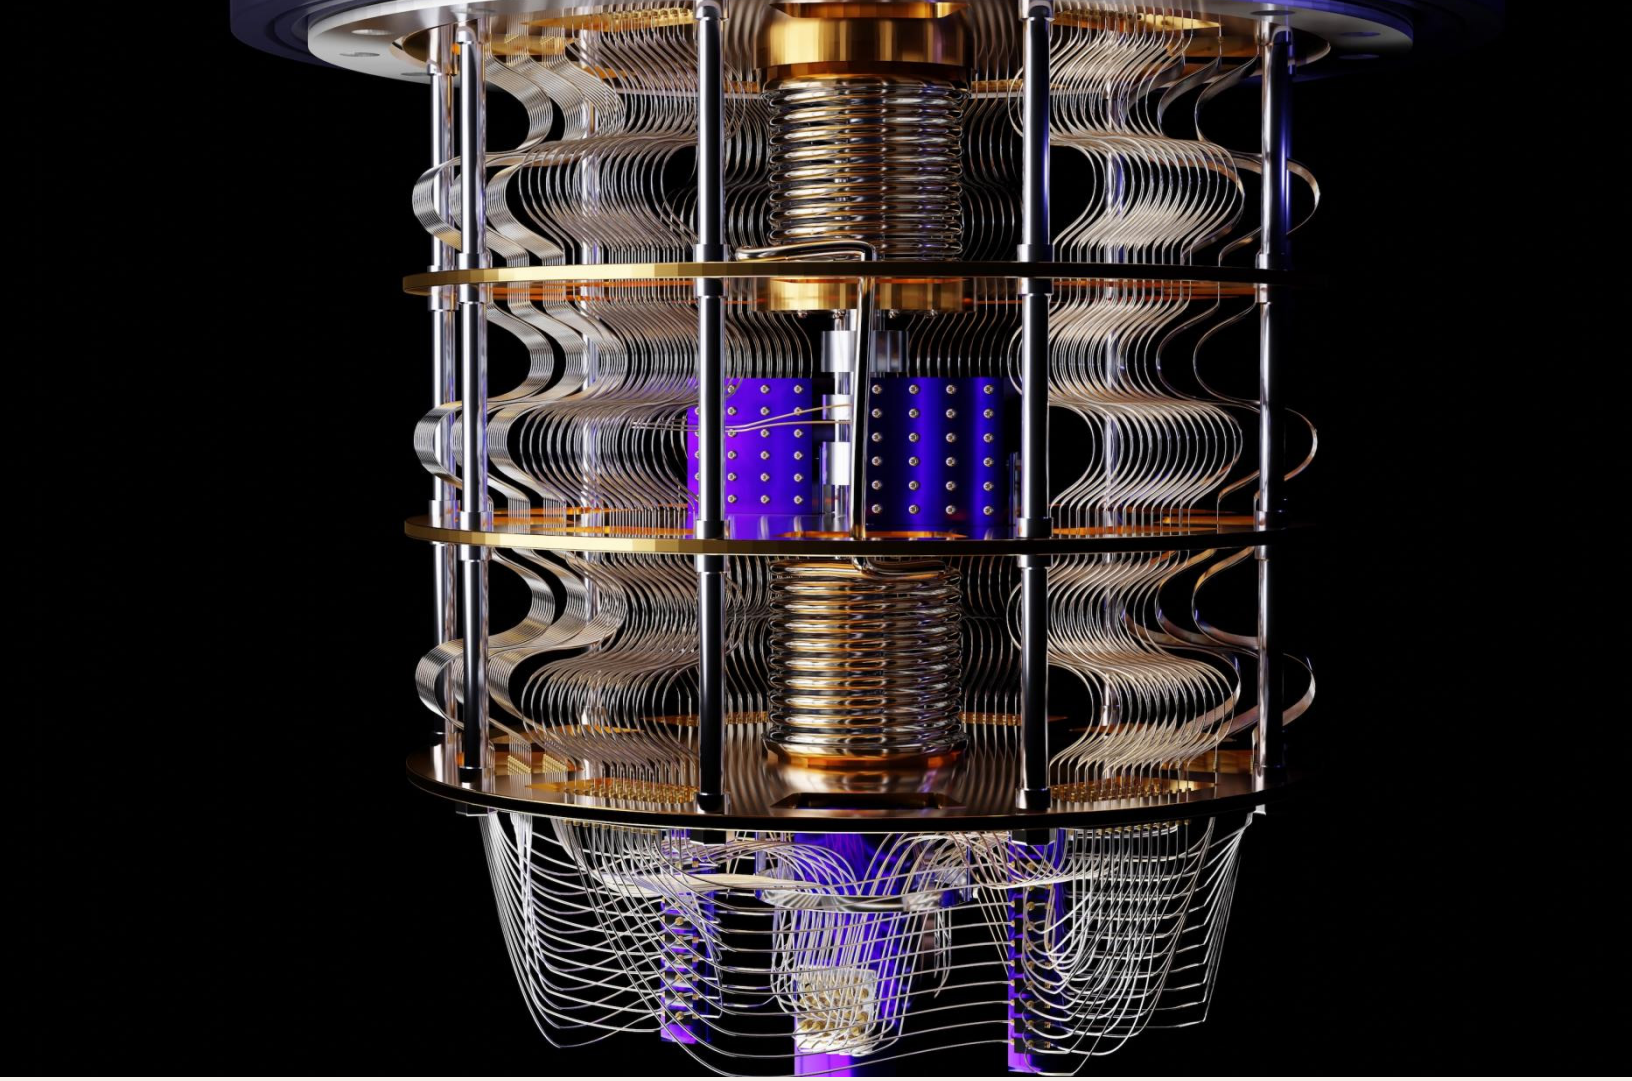


Zhijian Mao
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon

Decarbonization via Pulsed CO_2 Reduction Electrolyzers

Electrochemical CO_2 reduction can convert renewable electricity directly into ethanol and other multi-carbon (C_2+) products, but coupling an intermittent power source to a continuous chemical process remains a major engineering barrier. Pulsed electrolysis — alternating short current pulses with rest periods — has been proposed to both accommodate this intermittency and steer product selectivity away from the competing hydrogen evolution reaction and toward valuable C_2+ products. In a membrane electrode assembly built around a copper catalyst and a Sustainion anion-exchange membrane, we varied the pulse current density, pulse duration, and rest duration to probe how each parameter shapes the reaction environment at the catalyst surface and, in turn, the resulting product distribution. Across a systematic series of batches, gas products (H_2, CO, ethylene) and liquid products (ethanol, acetate, formate) were quantified by micro-GC and NMR after each run, benchmarked against a constant-current (DC) baseline, and tracked across sequential batches to assess reproducibility and membrane/catalyst degradation over extended operation. In parallel, we also tested an alternative membrane (Versogen) under the same protocol, investigating the mechanisms behind its failure and the membrane–catalyst interactions contributing to it. The resulting map of faradaic efficiency and selectivity across the pulsing parameter space provides a data-driven basis for operating renewable-powered CO_2 electrolyzers under variable power while maintaining high, reliable ethanol selectivity.





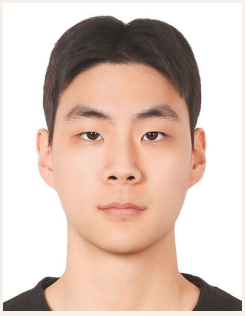
Quantum Science & Technologies

Quantum science will revolutionize technologies across multiple domains, from computing and communications, to sensors and secure data transmission, to global logistics and infrastructure security. However, the journey to realizing quantum's full promise has been hindered by a fragmented landscape across academia, industry, and government.

At NYU, we are changing this dynamic. We launched the NYU Quantum Institute to build a uniquely cross-sector, cross-disciplinary ecosystem of researchers, industry partners, and civic leaders to drive quantum innovation, education, workforce development, and real-world impact.

And we're already leading the charge with industry firsts: first quantum communication through existing telecom lines, the first academic cleanroom in Brooklyn, the first multi-node quantum network, and the first Thermal Laser Epitaxy system in the U.S. for advanced quantum materials research.

Quantum Science & Technologies



Sungyun Sohn
Mechanical Engineering, B.S.
Spring 2027
NYU Abu Dhabi

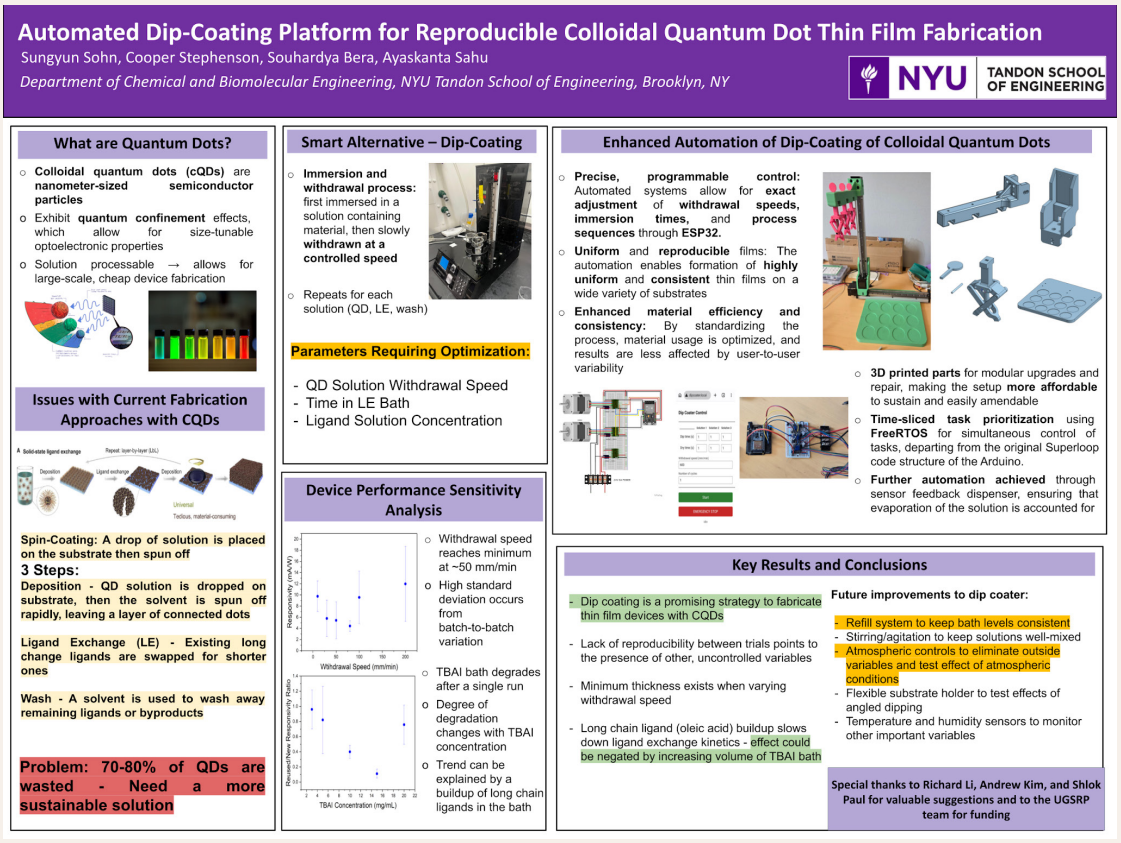


Cooper Stephenson
Chemical and Biomolecular Engineering, B.S.
Spring 2027
NYU Tandon

An Enhanced, Cost-Effective Automated Dip-Coating System for Quantum Dot Thin-Film Deposition

Dip coating is a simple yet effective method for depositing thin films. The process involves immersing a substrate material in a liquid solution and removing it to dry, consequently forming a thin film of solute following the evaporation of the solvent. Though the process itself is extensively researched, we explore its application to quantum dot film deposition. There are many parameters that affect the thickness and the uniformity of the film, which are critical factors that determine the quality of the products in optoelectronics applications. To test the influence of such variables like withdrawal speed, an automatic system that allows for precise control and incremental adjustments is needed. A commercially available dip coating machine capable of such control often ranges in price from $1500 to over $6000. An affordable alternative under $300 was proposed by the NYU Hybrid Nanomaterials Lab, which involves control of a linear rail at constant speeds using an Arduino. However, it faces stuttering issues at lower speeds, leaves no room for sensor feedback control, and requires human intervention in between cycles for solution height control.

This project proposes an enhanced version of the setup, using an ESP32 for faster speeds, FreeRTOS task management and wireless control through WiFi, along with additional upgrades such as a 3D printed substrate holder and solution dispenser to minimize human intervention during the deposition process. The departure from the Arduino superloop code structure allows the setup to concurrently execute tasks and incorporate sensor feedback control, which was not possible in the previous setup. Moreover, its barebone dual-axis setup leads to minimal solution perturbation, and leaves room for modular upgrades in the future.



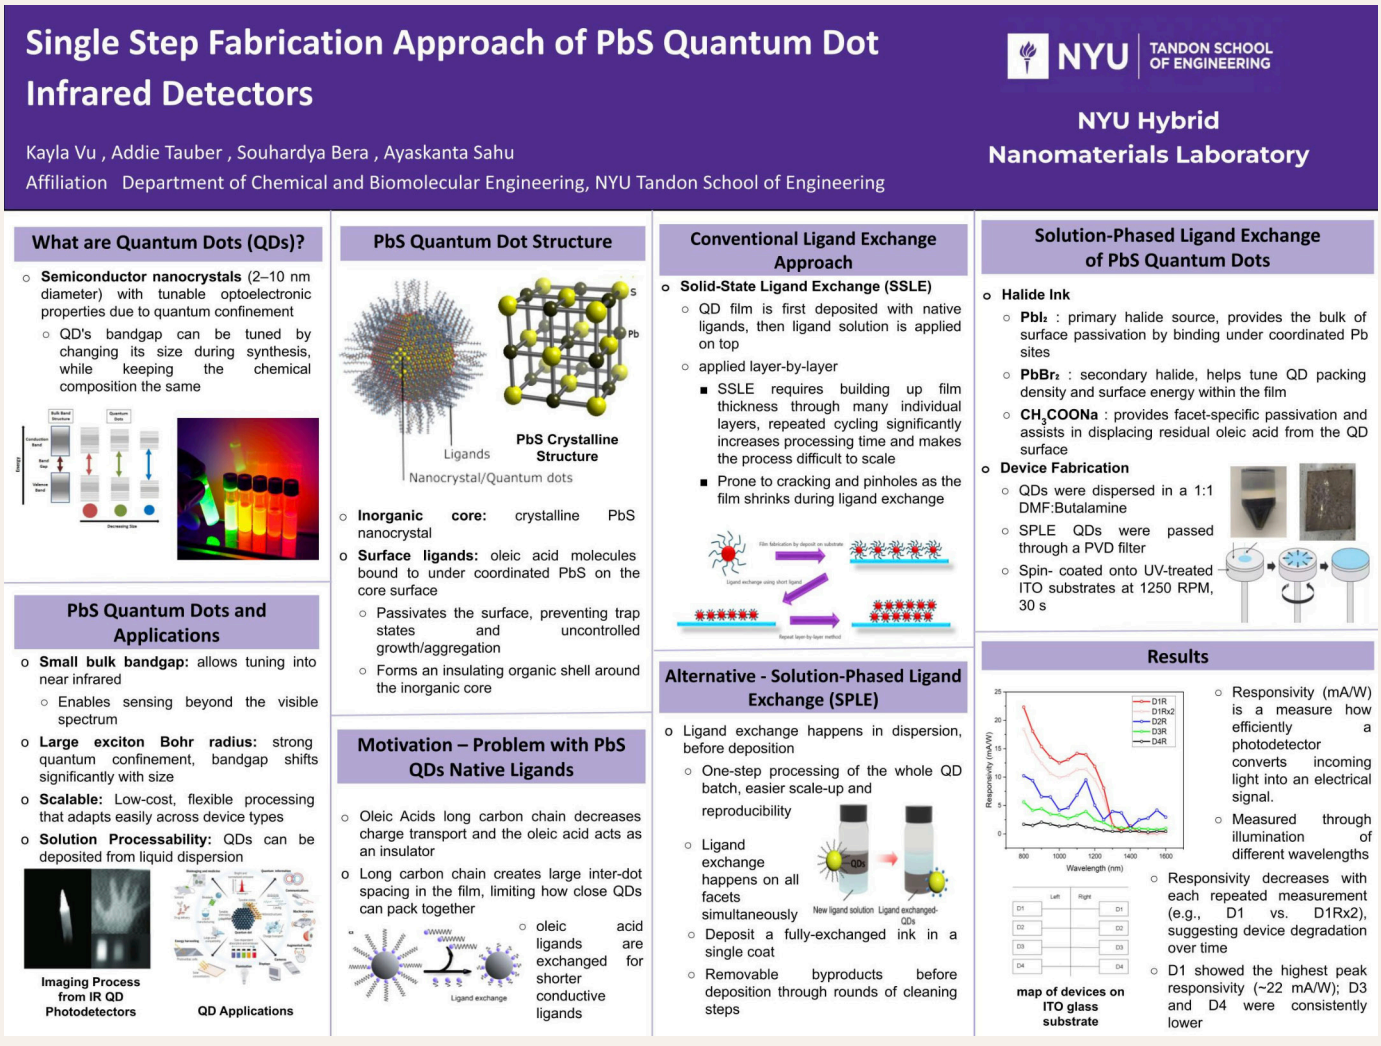
Quantum Science & Technologies



Kayla Vu
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon

Optimizing Solution-Phase Ligand Exchange for Enhanced Film Quality and Responsivity in PbS Colloidal Quantum Dot Photodetectors

Quantum dots (QDs) are semiconductor nanocrystals that possess size-tunable optical and electronic properties. PbS QDs specifically are synthesized with long oleic acid ligands, which prevent aggregation and provide high colloidal stability. However, the long carbon chain decreases charge transport and the oleic acid acts as an insulator. This project focused on solution-phase ligand exchange (SPLE) to improve film quality. A halide ink formulation combining PbI_2 and $PbBr_2$ was used to exchange oleic acid ligands for iodide and bromide. Additionally, NaOAc was incorporated for facet-specific passivation and assisting displacement of oleic acid. The resulting QDs were filtered and spin coated. Varying concentrations of PbI_2, $PbBr_2$, and NaOAc along with spin parameters and substrate cleaning procedures were tested to optimize film quality. Device performance was measured under illumination to capture responsivity measurements. Results showed that utilizing 0.150 M PbI_2, 0.03 M $PbBr_2$, 0.065 M NaOAc, UV treating ITO substrate, and spin coating the QDs at 1250 RPM for 30 seconds improved responsivity and consistency between devices. These findings demonstrate that precise control over halide ink composition, passivating agent concentration, substrate preparation, and deposition parameters is critical to producing reliable, high-performing PbS QD photodetector films. This work establishes an optimized SPLE protocol that improves reproducibility, supporting continued development of solution-processed infrared photodetectors.



Quantum Science & Technologies



Koki Furusho
Electrical Engineering and Computer Engineering (dual major), B.S.
Spring 2028
NYU Tandon



Sophia Wang
Major
Summer 2026
NYU Tandon

Kidney Stone Detection via Quantum Support Vector Machine on Real Quantum Hardware

Quantum computing has been attracting attention, including in the medical field. Especially for real-time surgery, a higher-dimensional expression and potential computational speedup would improve our current methods, which highly depend on humans. Here, we aim to evaluate whether Quantum Support Vector Machine (QSVM) can achieve clinically meaningful accuracy for kidney stone detection using real patient data, both on classical simulators and real IBM quantum hardware. Image features are extracted using ResNet-18, compressed via principal component analysis, and encoded onto qubits with a ZZ feature map before being passed through the QSVM's quantum kernel. While many research papers largely rely on simulation alone, we validate our pipeline on IBM's superconducting quantum processors, directly capturing the effects of hardware noise. Benchmarking against a conventional classical machine learning baseline across sample sizes ranging from 100 to 1000, we found that both accuracy and recall improved with sample size for both approaches. QSVM reached an accuracy of 0.70 for kidney stone detection, approaching the classical baseline of 0.77. This work serves as an early proof of concept for QSVM in real-time surgical image classification, highlighting both its near-term feasibility and the remaining hardware limitations toward clinical deployment.

NYU TANDON SCHOOL OF ENGINEERING

Kidney Stone Detection via Quantum Support Vector Machine on Quantum Hardware

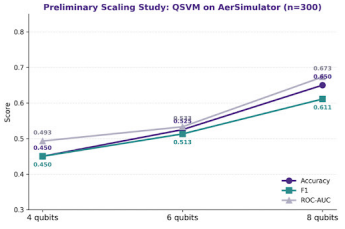
Koki Furusho[1,2], Sophia Wang[3], Rui Li[4]
[1]Electrical and Computer Engineering [2]Applied Physics [3] Computer Science and Engineering
[4]Mechanical and Aerospace Engineering, New York University

Abstract

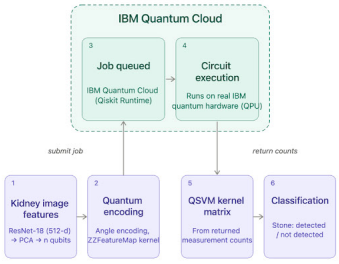
- Evaluated whether a Quantum Support Vector Machine (QSVM), tested on real IBM quantum hardware, can achieve clinically meaningful accuracy for kidney stone detection
- **QSVM reached 0.70 accuracy**, approaching the classical baseline of 0.77
- Demonstrates near-term feasibility of quantum ML for real-time surgical support, despite current hardware limitations

Introduction

- Real-time integration of computer vision requires a faster and more accurate algorithm
- Prior research showed that quantum algorithms might be solutions by running them on simulators [1,2,3,4]
- QSVM is expected to have a exponential speed up for classification tasks that take polynomial time for classical computers[5]

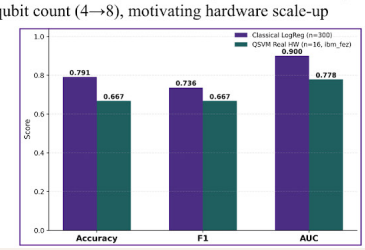


Algorithm



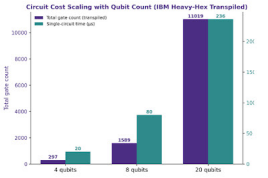
Data/Observations

- QSVM on real IBM hardware approached classical performance (F1 0.667 vs. 0.736) at n=16 vs. n=300
- Simulated accuracy/F1/AUC improve monotonically with qubit count (4→8), motivating hardware scale-up



Conclusion

- QSVM achieved competitive F1 vs. classical baselines at small scale, including **a successful real IBM hardware**
- Classical models still lead on accuracy/AUC; whether this quantum kernel can be classically simulated remains an open theoretical CS question
- Bottleneck: **$O(N^2)$ circuit cost** and limited QPU access restrict real-hardware sample size
- Future work: scale up real-hardware samples via **IBM collaboration**



Works Cited

[1] L. Wei et al., "Quantum machine learning in medical image analysis: A survey", Neurocomputing. **525,** 42–53 (2023).
[2] Sar, Ayan, et al. "Quantum Machine Learning in Medical Image Analysis: From Diagnostics to Surgery Planning." IEEE Access (2026).
[3] Matondo-Mvula, Nadine, and Khaled Elleithy. "Advances in quantum medical image analysis using machine learning: Current status and future directions." 2023 IEEE International conference on quantum computing and engineering (QCE). Vol. 1. IEEE, 2023.
[4] Yan, Fei, et al. "Review of medical image processing using quantum-enabled algorithms." Artificial Intelligence Review 57.11 (2024): 300.
[5] Rebentrost, Patrick, et al. "Quantum Support Vector Machine for Big Data Classification." Physical Review Letters, vol. 113, no. 13, Sept. 2014, p. 130503. American Physical Society, doi:10.1103/PhysRevLett.113.130503.

Acknowledgements

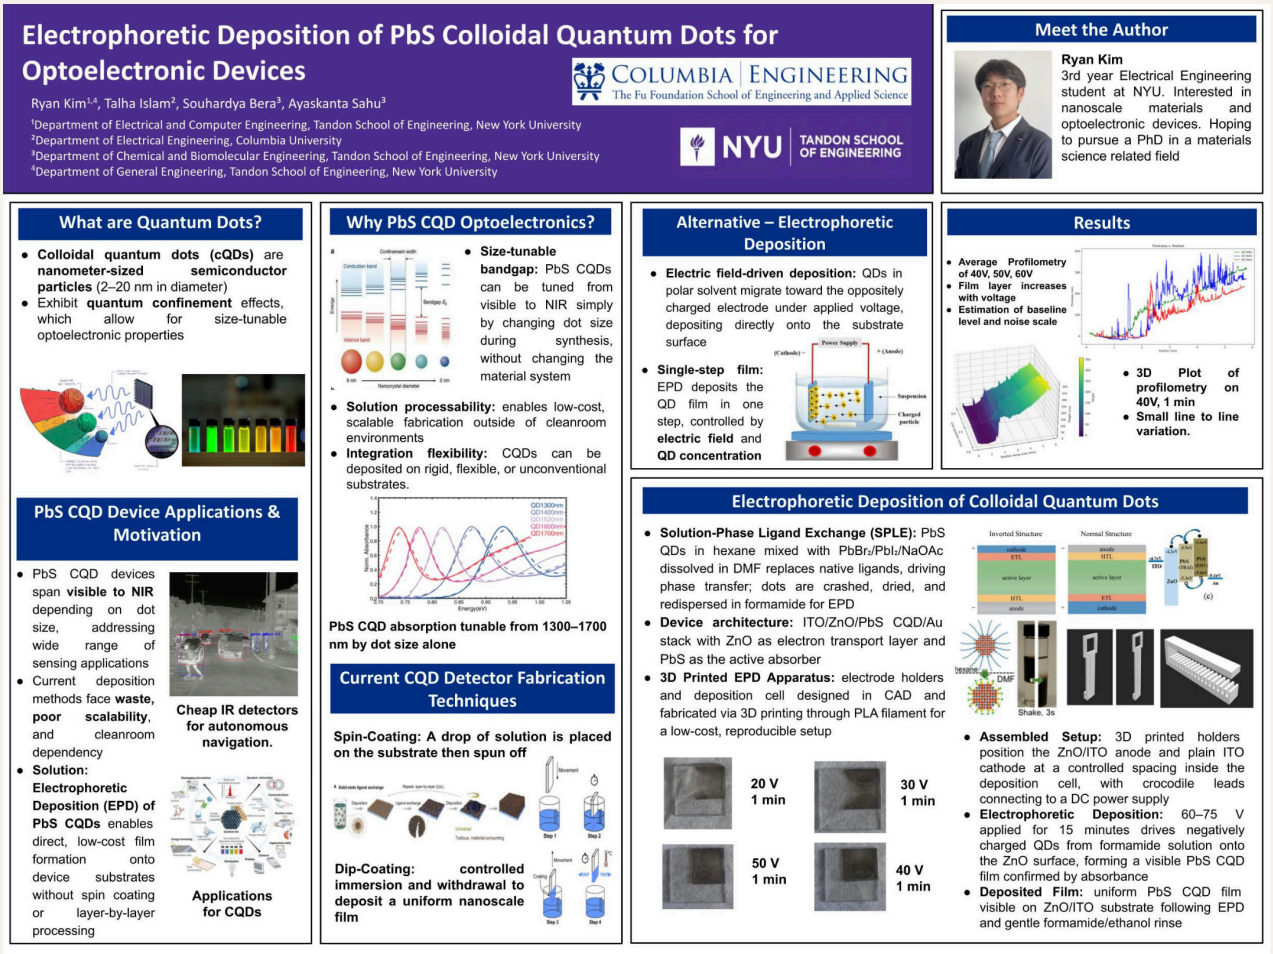
We thank E.E. Sanz, A Greggo and J. Tokuda for their insights and discussions.
This work is supported by UGSRP at NYU Tandon.

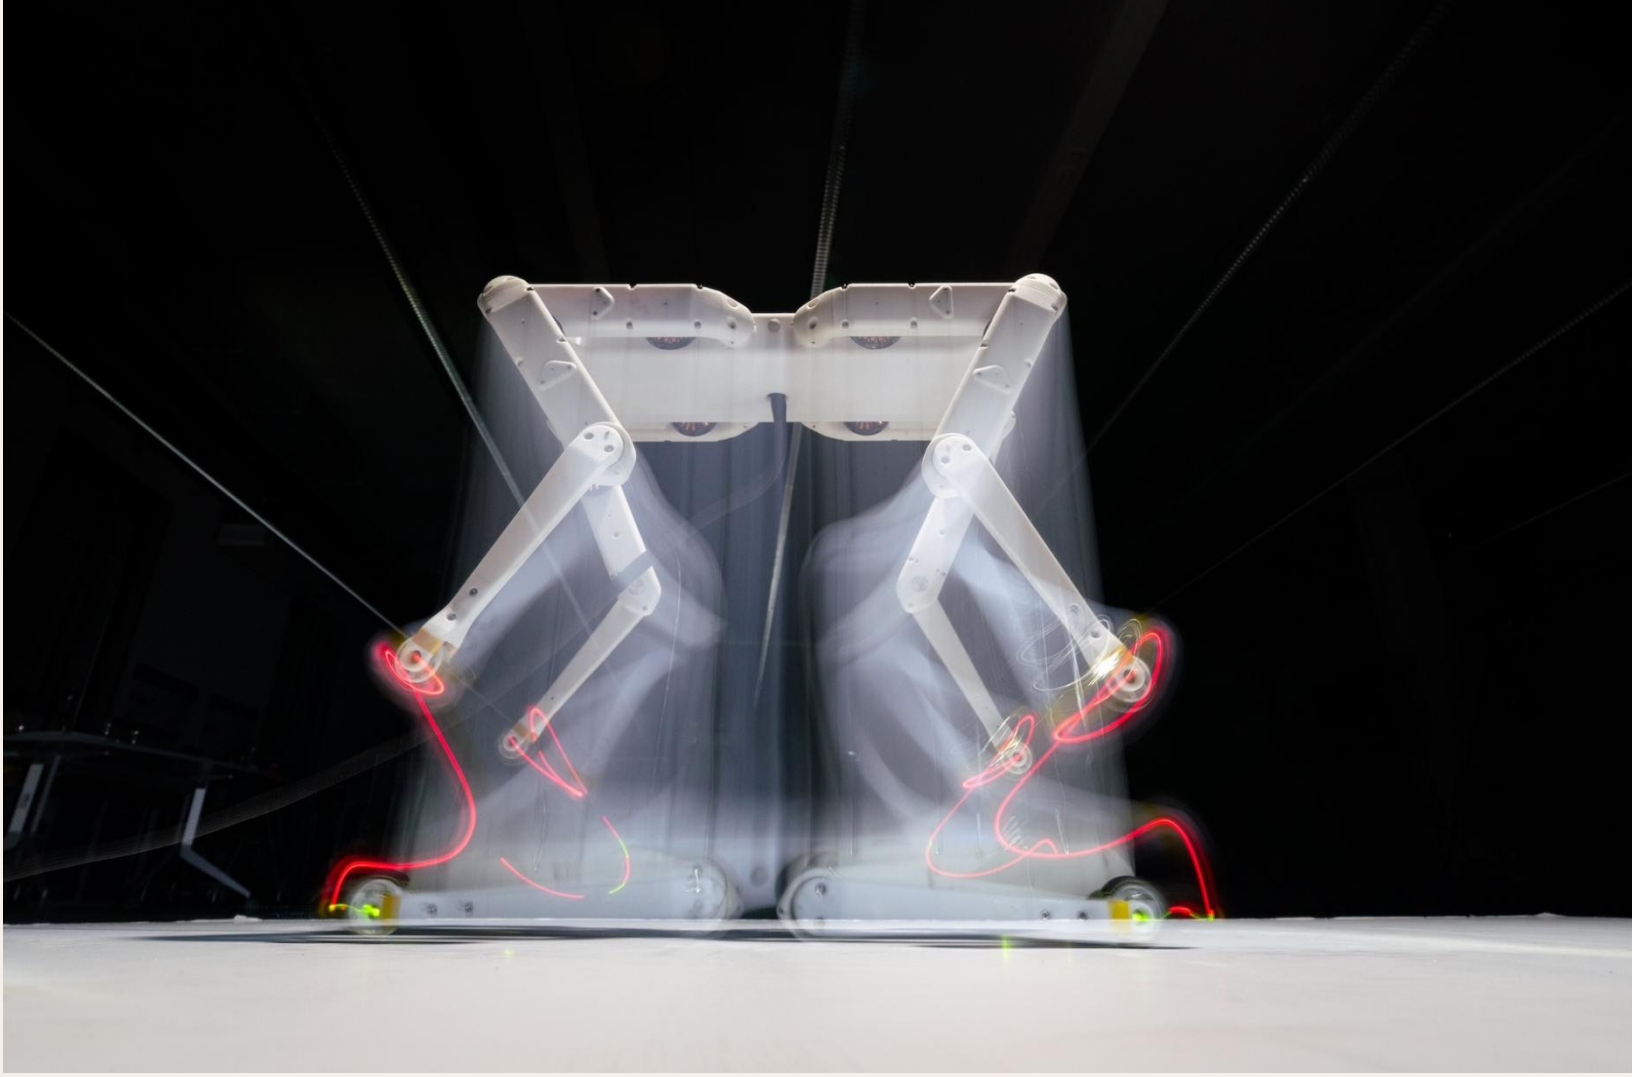


Ryan Kim
Chemical and Biomolecular Engineering, B.S.
Spring 2027
NYU Tandon

Low-Cost Electrophoretic Deposition of PbS Quantum Dot Films for Scalable Optoelectronics

Solution-processed optoelectronics have positioned colloidal quantum dots (CQDs) as promising candidates for next-generation solar cells and photodetectors, owing to their size-tunable bandgap and compatibility with low-cost, scalable manufacturing. Yet film deposition for CQD devices remains largely confined to spin-coating and layer-by-layer methods, both of which are time-consuming, material-wasteful, and difficult to scale. Electrophoretic deposition (EPD) offers a potentially simpler and more scalable alternative, but it has seen limited demonstration across diverse QD material systems and device architectures, particularly with solution-phase ligand exchanged (SPLE) inks. This work demonstrates EPD as a viable film deposition route for PbS CQDs on ZnO/ITO substrates, using a low-cost, 3D-printed deposition apparatus. PbS QDs synthesized via hot injection were subjected to SPLE with a mixed ligand system of $PbBr_2$, PbI_2, and sodium acetate to displace the native oleic acid capping and render the dots charged and dispersible in formamide for EPD. Deposition parameters were systematically varied, and 75 V applied for 30 minutes was identified as the optimal condition, producing a visibly uniform film with confirmed absorbance and measurable device responsivity. Higher voltages (100–200 V) instead resulted in non-uniform deposition and QD aggregation near the electrode contact point. These results demonstrate that EPD of SPLE PbS CQDs onto ZnO/ITO is achievable with minimal infrastructure, and that deposition quality is strongly governed by applied voltage. This work broadens the material and substrate combinations shown to be compatible with EPD, establishing a foundation for low-cost, scalable CQD film fabrication outside conventional cleanroom or glovebox environments.





Robotics & Embodied Intelligence

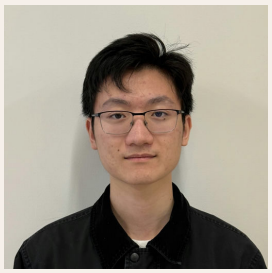
Robotics and Artificial Intelligence (AI) are both expansive fields far from reaching their full potential. Here at Tandon, we believe the most promising future lies at the intersection of the two with Embodied Intelligence where AI and wireless innovation merges with the physical world to create a new generation of smart, autonomous systems.

A technological revolution is approaching as AI innovations enable robots, self-driving vehicles, and smart cities to sense, learn, and act intelligently within their environments. By uniting robotics, machine learning, and computer vision, we are solving critical challenges in manufacturing, transportation, healthcare, and infrastructure while preparing and reshaping society for a more intelligent future.

Robotics & Embodied Intelligence



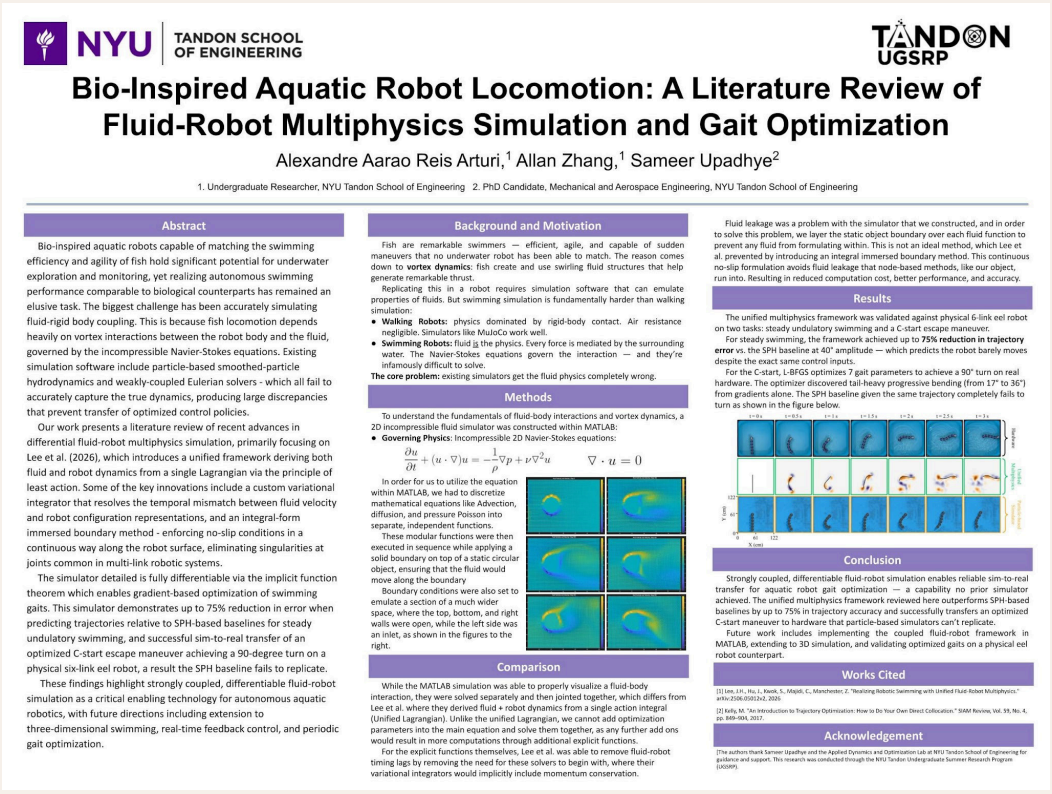
Alexandre Aarao Reis Arturi
Mechanical Engineering, B.S.
Winter 2028
University of Michigan



Allan Zhang
Mechanical Engineering, B.S.
Spring 2028
NYU Tandon

Bio-Inspired Aquatic Robot Locomotion: A Literature Review of Fluid-Robot Multiphysics Simulation and Gait Optimization

Bio-inspired aquatic robots capable of matching the swimming efficiency and agility of fish hold significant potential for underwater exploration and environmental monitoring, yet realizing autonomous swimming performance comparable to biological counterparts has remained an elusive task. The biggest challenge has been accurately simulating fluid-rigid body coupling. This is because fish locomotion depends heavily on vortex interactions between the robot body and the fluid, governed by the incompressible Navier-Stokes equations. Existing simulation software include particle-based smoothed-particle hydrodynamics and weakly-coupled Eulerian solvers - which all fail to accurately capture the true dynamics, producing large discrepancies that prevent transfer of optimized control policies. Our work presents a literature review of recent advances in differentiable fluid-robot multiphysics simulation, primarily focusing on Lee et al. (2026), wich introduces a unified framework deriving both fluid and robot dynamics from a single Lagrangian via the principle of least action. Some of the key innovations include a custom variational integrator that resolves the temporal mismatch between fluid velocity and robot configuration representations, and an integral-form immersed boundary method - enforcing no-slip conditions in a continuous way along the robot surface, eliminating singularities at joints common in multi-link robotic systems. The simulator detailed is fully differentiable via the implicit function theorem which enables gradient-based optimization of swimming gaits. This simulator demonstrates up to 75% reduction in error when predicting trajectories relative to SPH-based baselines for steady undulatory swimming, and successful sim-to-real transfer of an optimized C-start escape maneuver achieving a 90-degree turn on a physical six-link eel robot, a result the SPH baseline fails to replicate. These findings highlight strongly coupled, differentiable fluid-robot simulation as a critical enabling technology for autonomous aquatic robotics, with future directions including extension to three-dimensional swimming, real-time feedback control, and periodic gait optimization.



NYU TANDON SCHOOL OF ENGINEERING

TANDON UGSRP

Bio-Inspired Aquatic Robot Locomotion: A Literature Review of Fluid-Robot Multiphysics Simulation and Gait Optimization

Alexandre Aarao Reis Arturi,[1] Allan Zhang,[1] Sameer Upadhye[2]

1. Undergraduate Researcher, NYU Tandon School of Engineering 2. PhD Candidate, Mechanical and Aerospace Engineering, NYU Tandon School of Engineering

Abstract

Bio-inspired aquatic robots capable of matching the swimming efficiency and agility of fish hold significant potential for underwater exploration and monitoring, yet realizing autonomous swimming performance comparable to biological counterparts has remained an elusive task. The biggest challenge has been accurately simulating fluid-rigid body coupling. This is because fish locomotion depends heavily on vortex interactions between the robot body and the fluid, governed by the incompressible Navier-Stokes equations. Existing simulation software include particle-based smoothed-particle hydrodynamics and weakly-coupled Eulerian solvers - which all fail to accurately capture the true dynamics, producing large discrepancies that prevent transfer of optimized control policies.

Our work presents a literature review of recent advances in differential fluid-robot multiphysics simulation, primarily focusing on Lee et al. (2026), which introduces a unified framework deriving both fluid and robot dynamics from a single Lagrangian via the principle of least action. Some of the key innovations include a custom variational integrator that resolves the temporal mismatch between fluid velocity and robot configuration representations, and an integral-form immersed boundary method - enforcing no-slip conditions in a continuous way along the robot surface, eliminating singularities at joints common in multi-link robotic systems.

The simulator detailed is fully differentiable via the implicit function theorem which enables gradient-based optimization of swimming gaits. This simulator demonstrates up to 75% reduction in error when predicting trajectories relative to SPH-based baselines for steady undulatory swimming, and successful sim-to-real transfer of an optimized C-start escape maneuver achieving a 90-degree turn on a physical six-link eel robot, a result the SPH baseline fails to replicate.

These findings highlight strongly coupled, differentiable fluid-robot simulation as a critical enabling technology for autonomous aquatic robotics, with future directions including extension to three-dimensional swimming, real-time feedback control, and periodic gait optimization.

Background and Motivation

Fish are remarkable swimmers — efficient, agile, and capable of sudden maneuvers that no underwater robot has been able to match. The reason comes down to **vortex dynamics**: fish create and use swirling fluid structures that help generate remarkable thrust.

Replicating this in a robot requires simulation software that can emulate properties of fluids. But swimming simulation is fundamentally harder than walking simulation:

- **Walking Robots:** physics dominated by rigid-body contact. Air resistance negligible. Simulators like MuJoCo work well.
- **Swimming Robots:** fluid is the physics. Every force is mediated by the surrounding water. The Navier-Stokes equations govern the interaction — and they're infamously difficult to solve.

The core problem: existing simulators get the fluid physics completely wrong.

Methods

To understand the fundamentals of fluid-body interactions and vortex dynamics, a 2D incompressible fluid simulator was constructed within MATLAB:

- **Governing Physics**: Incompressible 2D Navier-Stokes equations:

$$\frac{\partial u}{\partial t} + (u \cdot \nabla)u = -\frac{1}{\rho}\nabla p + \nu\nabla^2 u \qquad \nabla \cdot u = 0$$

In order for us to utilize the equation within MATLAB, we had to discretize mathematical equations like Advection, diffusion, and pressure Poisson into separate, independent functions.

These modular functions were then executed in sequence while applying a solid boundary on top of a static circular object, ensuring that the fluid would move along the boundary

Boundary conditions were also set to emulate a section of a much wider space, where the top, bottom, and right walls were open, while the left side was an inlet, as shown in the figures to the right.

Comparison

While the MATLAB simulation was able to properly visualize a fluid-body interaction, they were solved separately and then jointed together, which differs from Lee et al. where they derived fluid + robot dynamics from a single action integral (Unified Lagrangian). Unlike the unified Lagrangian, we cannot add optimization parameters into the main equation and solve them together, as any further add ons would result in more computations through additional explicit functions.

For the explicit functions themselves, Lee et al. was able to remove fluid-robot timing lags by removing the need for these solvers to begin with, where their variational integrators would implicitly include momentum conservation.

Fluid leakage was a problem with the simulator that we constructed, and in order to solve this problem, we layer the static object boundary over each fluid function to prevent any fluid from formulating within. This is not an ideal method, which Lee et al. prevented by introducing an integral immersed boundary method. This continuous no-slip formulation avoids fluid leakage that node-based methods, like our object, run into. Resulting in reduced computation cost, better performance, and accuracy.

Results

The unified multiphysics framework was validated against physical 6-link eel robot on two tasks: steady undulatory swimming and a C-start escape maneuver.

For steady swimming, the framework achieved up to **75% reduction in trajectory error** vs. the SPH baseline at 40° amplitude — which predicts the robot barely moves despite the exact same control inputs.

For the C-start, L-BFGS optimizes 7 gait parameters to achieve a 90° turn on real hardware. The optimizer discovered tail-heavy progressive bending (from 17° to 36°) from gradients alone. The SPH baseline given the same trajectory completely fails to turn as shown in the figure below.

Conclusion

Strongly coupled, differentiable fluid-robot simulation enables reliable sim-to-real transfer for aquatic robot gait optimization — a capability no prior simulator achieved. The unified multiphysics framework reviewed here outperforms SPH-based baselines by up to 75% in trajectory accuracy and successfully transfers an optimized C-start maneuver to hardware that particle-based simulators can't replicate.

Future work includes implementing the coupled fluid-robot framework in MATLAB, extending to 3D simulation, and validating optimized gaits on a physical eel robot counterpart.

Works Cited

[1] Lee, J.H., Hu, J., Kwok, S., Majidi, C., Manchester, Z. "Realizing Robotic Swimming with Unified Fluid-Robot Multiphysics." arXiv:2506.05012v2, 2026.

[2] Kelly, M. "An Introduction to Trajectory Optimization: How to Do Your Own Direct Collocation." SIAM Review, Vol. 59, No. 4, pp. 849–904, 2017.

Acknowledgement

[The authors thank Sameer Upadhye and the Applied Dynamics and Optimization Lab at NYU Tandon School of Engineering for guidance and support. This research was conducted through the NYU Tandon Undergraduate Summer Research Program (UGSRP).

Robotics & Embodied Intelligence



Aman Sunesh
Computer Engineering, B.S.
Spring 2027
NYU Abu Dhabi

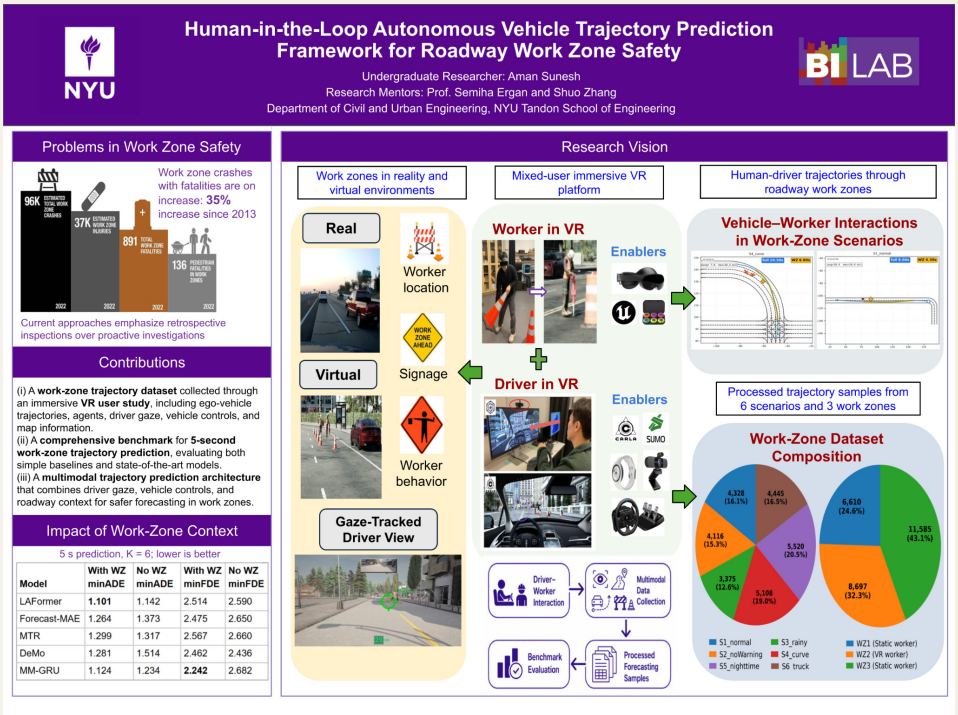
Human-in-the-Loop Autonomous Vehicle Trajectory Prediction Framework for Roadway Work Zone Safety

Trajectory prediction is a critical component of autonomous driving systems because it supports safe and effective motion planning. However, ensuring autonomous vehicle (AV) safety in safety-critical scenarios remains a significant challenge due to the scarcity of real-world data required to train robust data-driven models. Roadway work zones are common but highly dynamic driving environments. They often include cones, lane shifts, warning signs, equipment, and workers moving near traffic, all of which can make driver behavior harder to predict. Yet high-quality data for studying these work-zone scenarios remains limited.

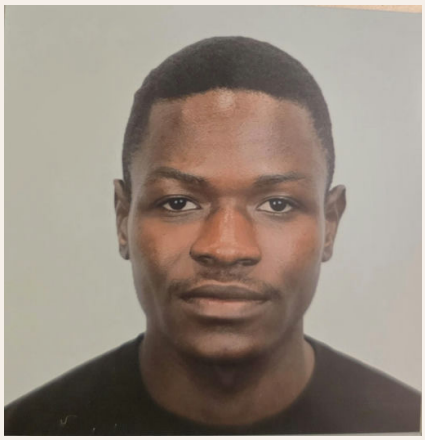
Building upon our previous work, which used an interactive virtual-reality (VR) framework to capture the trajectories and interactions of human drivers and workers in a controlled user study, this research addresses the critical gap in AV trajectory prediction for construction-zone environments. We introduce and release a novel human-derived trajectory dataset designed for this safety-critical domain. The dataset includes vehicle motion, surrounding agents, work-zone geometry, warning signs, control inputs, gaze information, and map context. We cleaned and processed the data from two laps of the study into model-ready trajectory samples with consistent train, validation, and test splits by participant.

Furthermore, we establish a comprehensive evaluation benchmark for predicting future vehicle motion over a five-second horizon in these complex environments. Our benchmark evaluates a range of predictive models, from foundational sequence models such as GRU, LSTM, and Transformer baselines to more recent trajectory prediction architectures, including LAformer, MTR, and RouteFormer. Initial results show that work-zone context and vehicle-control information improve prediction performance, while richer map and lane features may require more specialized model architectures to be fully useful.

Ultimately, this study aims to advance safe autonomous navigation in work zones by providing a benchmarked foundation for improving the reliability of AV motion planning and developing more accurate trajectory-prediction models in complex work-zone environments.



Robotics & Embodied Intelligence



Demarce Williams
Electrical Engineering, B.S.
Spring 2027
NYU Abu Dhabi

What Do Spectrum Prediction Models Learn? A Benchmark Dataset and Reproduction Study

Spectrum usage prediction has attracted growing interest from researchers, engineers, and industry due to its potential to reveal usage patterns and support resource allocation decisions. However, translating prediction models into reliable field deployments remains difficult, as evaluation methods often fail to align with the capabilities researchers claim their models possess. As a result, overly optimistic evaluations may yield models that perform well offline but fail in deployment, leading to inaccurate spectrum forecasts and poor downstream allocation decisions. In this work, we develop a best-practice framework for evaluating spectrum prediction models by identifying specific claims made in the literature, constructing datasets and evaluation procedures designed to test each claim directly, and reimplementing and reevaluating selected existing models using these benchmarks. This framework provides a more reliable measure of deployment readiness, helps identify model limitations, highlights areas for improvement, and clarifies whether spectrum prediction models are worth their training and deployment costs. Ultimately, more rigorous evaluation can lead to more robust models better suited for real-world deployment.

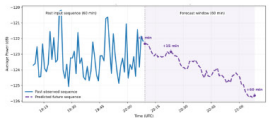
DEMARCE WILLIAMS
New York University Abu Dhabi
Advised by: Prof. Fraida Fund &
Prof. Shivendra Panwar
NYU Tandon School of Engineering

What Do Spectrum Prediction Models Learn?
A Benchmark Dataset and Reproduction Study

Simple baselines matched or outperformed complex neural models.

Spectrum forecasts could improve wireless networks

Predicting future activity could help wireless systems identify when frequencies will be available and share spectrum more efficiently.



Prior work promises good results, can it deliver?

Many models in the literature are evaluated using:

- one frequency band or location;
- short measurement periods;
- inconsistent evaluation procedures;
- limited baseline comparisons.

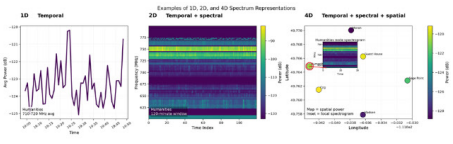
so it's not clear whether their findings would hold up in deployment.

We develop a new benchmark dataset and re-evaluate nine models from the literature

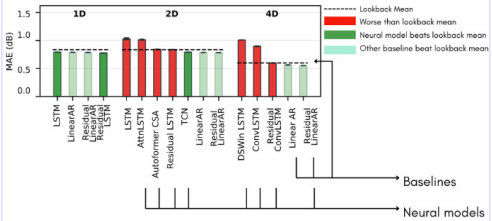
Testbeds	POWDER, COSMOS, ARA, AERPAW
Environments	Campus, urban, and rural
Other factors	Receivers, frequency bands, and day
Data	> 3,000 hours of data

Three representations isolate distinct structures

We represented the data at three levels to isolate the temporal, spectral, and spatial structure available to each model.

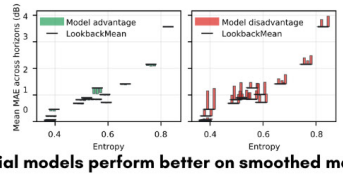


Core finding: Simple baselines beat or match complex neural models

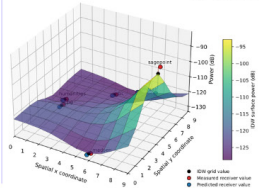


Neural Models didn't do better on more difficult bands

Entropy measures how "unpredictable" a frequency band is. Complex neural models didn't have an advantage on high-entropy bands



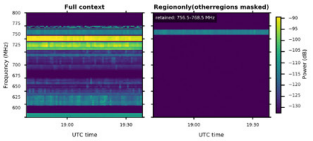
Spatial models perform better on smoothed maps than at real receivers



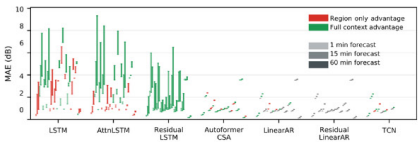
Lower map error does not prove spatial learning.

0.861 dB	Full IDW grid
1.134 dB	Real receiver IDW node
1.982 dB	Raw receiver node

Spectral information can help, but its benefit is architecture-dependent.



Full spectral context reduced ResidualLSTM MAE from 2.419 to 1.197 dB, but produced little improvement for the strongest models.



Simple models remained robust under transfer distribution shift

error > 20 dB
LSTM, AttentionLSTM, TCN

error < 0.1 dB
Lookback Mean, Residual LinearAR

Simple and residual models transferred reliably; AutoformerCSA was a robust neural exception, while TimeRAN did not outperform MOMENT.

Complexity alone is not evidence of meaningful learning

- Strong in-domain performance does not prove that a model learns transferable spectrum structure.
- Neural models should be tested against strong simple baselines.
- Claims about temporal, spectral, or spatial learning require controlled ablations.
- Evaluation should extend across times, bands, receivers, and testbeds.

References & Acknowledgements

D. S. Williams, F. Fund, S. S. Panwar, and L. Mangal, "What Do Spectrum Prediction Models Learn? A Benchmark Dataset and Reproduction Study," manuscript under review, 2026.

Acknowledgements: Conducted at NYU WIRELESS under Prof. Shivendra S. Panwar and Dr. Fraida Fund, using POWDER, COSMOS, ARA, AERPAW, and Chameleon.

NYU WIRELESS TANDON UGSRP

Robotics & Embodied Intelligence



He Tu
Data Science(AI concentration) and Mathematics , B.S.
Spring 2027
NYU Shanghai

Coverage-Aware Retrieval-Augmented Generation for Multifaceted Question Answering

While Large Language Models (LLMs) continue to push boundaries, they still struggle with hallucinated outputs, stale data, and untraceable reasoning. Retrieval-Augmented Generation (RAG) offers a promising solution by drawing from external databases. However, most existing RAG frameworks focus heavily on simple, factual QA. When tasked with nuanced, open-ended questions requiring thorough responses, they frequently provide incomplete answers. The bottleneck comes down to a fundamental objective mismatch in current retrieval design. Generating truly comprehensive answers requires high information coverage, retrieving a collection of documents that together span every critical angle of a query. Standard retrievers, however, rank documents independently. This creates severe nugget-level redundancy: top search results endlessly rehash the same primary aspect, completely ignoring other necessary dimensions. To overcome this limitation, we introduce a coverage-aware retrieval pipeline designed to maximize distinct information nuggets within a limited context window. Instead of static, isolated ranking, our framework uses a dynamic set-selection strategy that pairs a marginal-gain retriever and a distilled LLM coverage scorer with a greedy module that actively hunts for missing information gaps. Evaluated on the CoverageBench benchmark, our approach improves both retrieval relevance and coverage, showing promising potential to boost final answer quality.

Coverage-Aware Retrieval-Augmented Generation

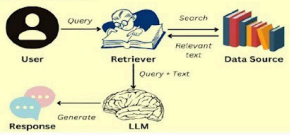
He Tu,[1] Aleena Bosco Kurumthottam,[2] Torsten Suel[2]

1. NYU Shanghai
2. NYU Tandon School of Engineering

NYU TANDON SCHOOL OF ENGINEERING

Abstract

Retrieval-Augmented Generation (RAG) improves LLMs with external knowledge, but standard retrieval ranks documents independently, often retrieving redundant evidence and missing key aspects of complex questions. We propose a coverage-aware retrieval pipeline that maximizes information coverage within a limited context window.



Background

Concept	Description
Hallucination	Unsupported information generated by LLMs.
Generation	Produces answers from retrieved context.
Standard Retrieval	Ranks documents independently by relevance.
RAG	Retrieves external knowledge before generation.
Grounding	Supports answers with retrieved evidence.
Multi-Faceted Questions	Require evidence from multiple aspects.
Relevance	Measures document-query matching.
Coverage	Measures distinct information covered.

Research Question

- Can retrieval optimize set-level coverage instead of document-level relevance?
- Does better retrieval coverage lead to more complete LLM-generated answers?

Framework Design

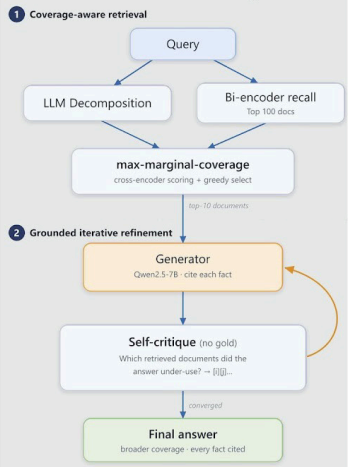


Figure 1: The coverage-aware iterative RAG pipeline

Data

- **Training:**
  Researchy Question dataset
- **Evaluation:**
  RAG24 dataset
  Qwen2.5-7B generator.

Result

Step1: Retrieval

Method	α-nDCG	StRec	P	nDCG
Baseline	0.625	0.744	0.754	0.784
+ coverage training	0.656	0.762	0.773	0.798
+ cross-enc & greedy	0.702	**0.781**	0.775	0.800
Gold-facet (UB)	**0.718**	**0.781**	0.771	**0.802**

Table 1: Retrieval coverage on RAG24 Retrieval coverage on RAG24. Each row adds one component over the previous.

- Coverage training consistently improves retrieval coverage.
- Our full method nearly reaches the gold-facet upper bound.

Step2: Generation Side:

Metric	A	B	B+It	O
StRec@k	0.791	0.799	0.799	**0.874**
Transfer	0.294	0.326	**0.376**	0.270
Ans-SR	0.239	0.266	**0.289**	0.242
V-strict	0.350	0.338	**0.406**	0.323
Sente_prec	0.715	**0.703**	0.652	0.627

Table 2: Downstream generation A = relevance-only retrieval; B = coverage-aware retrieval; B+Iter = grounded iterative refinement; O = oracle retrieval.

- Coverage-aware retrieval conveys more nuggets than oracle retrieval.
- Iterative refinement further improves coverage while maintaining faithfulness.

Limitation

- Evaluated on one benchmark only.
- Single-round refinement,self-detection quality limits performance.

Conclusion

- Set-level coverage retrieval beats per-document relevance — and conveys more nuggets downstream than even a ground-truth oracle.
- The remaining bottleneck is generation, not retrieval: more retrieved ≠ more conveyed.
- A grounded, self-detecting iterative loop (no gold) recovers coverage *while keeping faithfulness* → a deployable coverage-aware RAG.

Future Work

- Implement on more datasets to push answer-coverage gains to statistical significance.
- Use a larger mode for self-detection to close the remaining gap to the oracle.
- Multi-round loop: iterate the self-critique beyond one round for further cover age.

Works Cited

[1] Ju et al., "Search for Coverage: Learning Coverage-Aware Retrieval with Augmented Sub-Question Answerability," Proc. SIGIR, 2026.
[2] Samuel et al., "CoverageBench: Evaluating Information Coverage across Tasks and Domains," arXiv:2603.20034, 2026.
[3] Rosset et al., "Researchy Questions: A Dataset of Multi-Perspective, Decompositional Questions for LLM Web Agents," arXiv:2402.17896, 2024.
[4] Liu et al., "Lost in the Middle: How Language Models Use Long Contexts," arXiv:2307.03172, 2023.

Acknowledgement

This work was supported by the NYU Tandon School of Engineering. We thank the school for its funding, computational resources, and research support.

Robotics & Embodied Intelligence



Huiyi Yang
Computer Science, B.S.
Spring 2027
NYU Shanghai

Wanderland × Video Generation: Dynamics, Visual Quality, and Consistency for a Real-Scan Navigation Simulator

Real-scan simulators give embodied navigation metrically grounded environments, but pure reconstruction hits three ceilings: quality collapses off the capture trajectory (best 3DGS: 20.4 dB interpolation → 17.9 dB extrapolation), glass and thin structures fail intrinsically under LiDAR, and scans are captured people-free, so there are no dynamics at all. We keep geometry measured and use video generation only to paint what geometry cannot. A fixed colored LiDAR cloud is point-rendered at any query pose and decoded by a frozen Wan2.2-5B video model with a VACE conditioning branch and LoRA; because every view is conditioned on the same cloud, spatial consistency holds by construction rather than by learned memory. One model across 10 scenes reaches 24.5 dB on held-out validation, matching per-scene 3DGS while replacing ten reconstructions with one, and supports never-photographed viewpoints — lateral offsets, yaw sweeps, and goal-driven rollouts planned directly on the scan. For dynamics we separate behavior from appearance: deterministic proxy tracks with explicit z-testing against the LiDAR depth buffer make actor state exact and occlusion a property of the condition, leaving only appearance to generation. We pair held-out-photo metrics with ground-truth-free consistency checks available at any pose.

NYU TANDON SCHOOL OF ENGINEERING

Improving the Visual Quality of Real-World Scanned Scenes through Geometry-Aware Video Generation

Huiyi Yang

Background

Real-world scans preserve authentic geometry, object layouts, materials, lighting, and camera trajectories, reducing the gap between simulation and deployment.
However, scanning often produces missing regions, reflective-surface failures, boundary artifacts, noisy geometry, and poor novel-view rendering. Generative refinement may improve visual quality, but it can also hallucinate objects or alter scene structure.
Challenge: Enhance appearance while preserving geometric consistency.

Contribution & Limitation

We combine latent Gaussian memory with VACE conditioning to improve noisy scan renderings. Results outperform raw rendering, though 3DGS remains stronger near captured views.
Limitations & Future Work

- Weakly observed regions may contain hallucinated objects or geometry.
- Autoregressive generation can accumulate temporal errors.
- Future work will improve geometric constraints, temporal consistency, memory updates, and larger-scale evaluation.

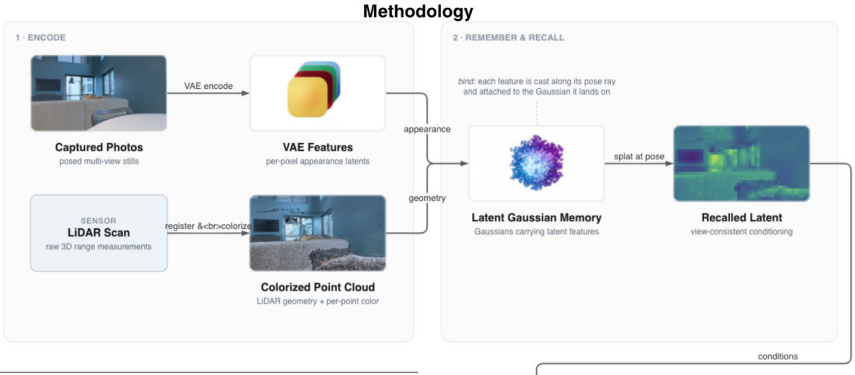
Works Cited

- Chen et al. (2026) introduce Robust Dreamer, which learns latent Gaussian memory from video for geometry-aware autoregressive generation.
- Wang et al. (2026) propose Image2Sim, a generative neural simulator constructed from posed RGB-D observations.
- de Lutio et al. (2026) present AttiFixer, an autoregressive diffusion framework for repairing and extending incomplete 3D reconstructions.
- Wu et al. (2025) propose DiFix3D+, which improves NeRF and 3DGS reconstructions using single-step diffusion refinement.

Acknowledgement

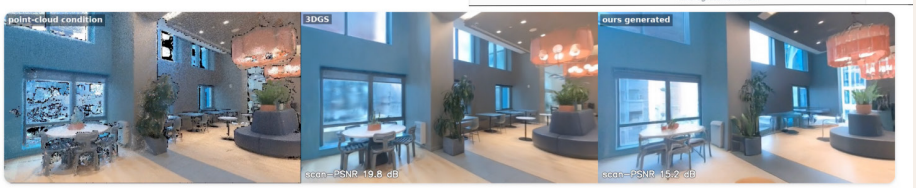
The authors gratefully acknowledge the NYU Tandon School of Engineering Office of Undergraduate Academics for its generous support of this project.

Methodology



Experiments & Results

Method	PSNR ↑	LPIPS ↓	Scan-PSNR ↑
Raw Render	14.1	0.894	–
3DGS	22.0	0.290	19.9
Ours	16.0	0.432	15.2



Robotics & Embodied Intelligence



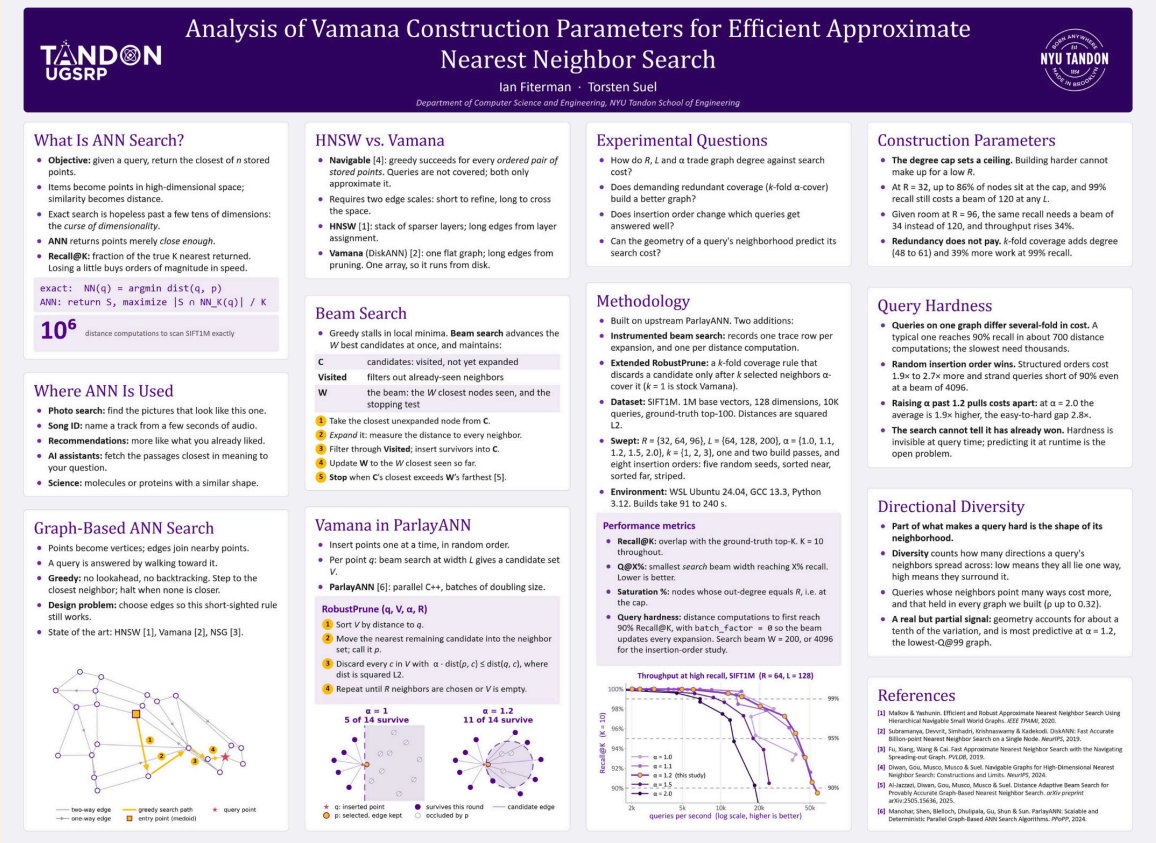
Ian Fiterman
NYU Tandon
Spring 2027
Computer Science, B.S.

Analysis of Vamana Construction Parameters for Efficient Approximate Nearest Neighbor Search

Modern systems represent images, documents, and other media as high-dimensional vectors and answer queries by retrieving those closest to a query. Scanning every stored vector is exact but too slow at scale, so production systems use approximate nearest neighbor (ANN) search, trading a little accuracy for speed. The strongest methods build a graph and walk it greedily toward the query. Vamana, the algorithm behind DiskANN, builds that graph by inserting points one at a time, using beam search to gather nearby candidates, then pruning so each node keeps few edges in varied directions. Its behavior depends on the maximum out-degree R, the construction beam width L, a pruning threshold α governing how readily redundant edges are dropped, and the insertion order. These are inherited as defaults with little published sensitivity analysis.

We instrumented ParlayANN's Vamana with per-distance-computation tracing, a graph serializer, insertion-order control, and a generalized pruning rule, then swept these parameters on SIFT1M, one million 128-dimensional vectors. R is the dominant parameter as setting it too low results in poor quality graphs despite higher construction effort. L helps only once R leaves headroom for the additional edges discovered with more computation during beam search. α has a narrow optimum near 1.2, the accepted default, since lower values strip the long edges that let search travel quickly, while higher values crowd nodes with near-duplicates that slow every step. Query difficulty was only partly shared across insertion-order variants, so it is neither entirely a property of the dataset nor of the graph, and more work is needed to explain the remaining variance.

Our main outcome is a clearer account of why Vamana's chosen settings work. They can be found by exhaustive sweeps, but understanding the mechanism lets us reason about them and tune with far fewer experiments.



Robotics & Embodied Intelligence



Junho So
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon



Michael Carcamo
Computer Engineering, B.S.
Spring 2028
NYU Tandon



Nazmus Sakib
Computer Science, B.S.
Spring 2028
NYU Tandon



Noah Mays-Smith
Computer Engineering, B.S.
Spring 2026
NYU Tandon



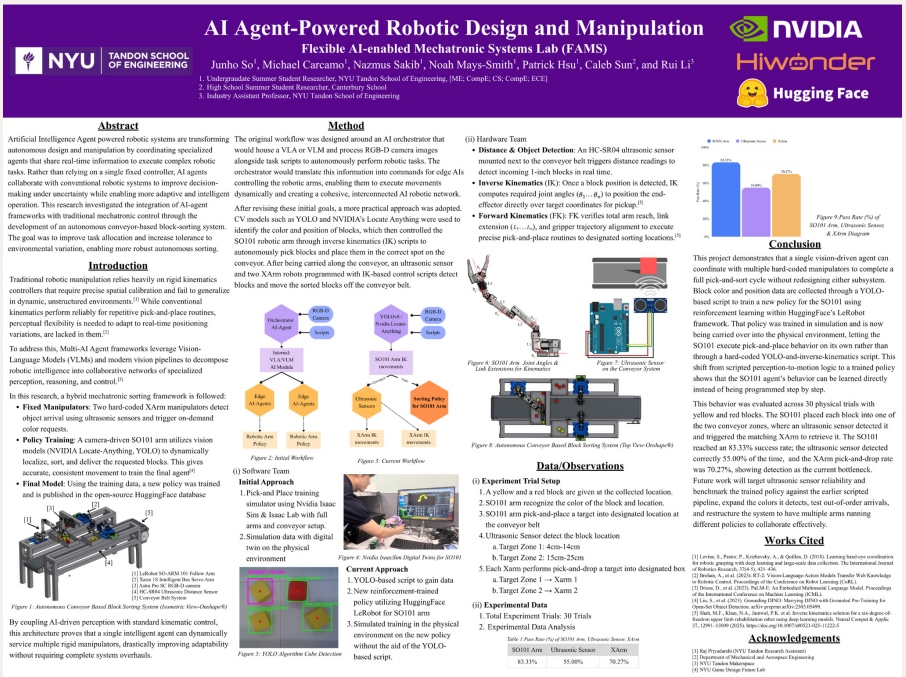
Patrick Hsu
Electrical and Computer Engineering, B.S.
Spring 2026
NYU Tandon

AI Agent-Powered Robotic Design and Manipulation

AI-agent-powered robotic systems advance autonomous design, fabrication, and manipulation by coordinating specialized AI models to co-reason, co-design, and co-execute complex robot behaviors. Instead of relying on standard controllers, AI agents share real-time information to adaptively select execution strategies under uncertainty, supporting both human interaction and in-silico co-design loops. This research explored the possibility of integrating AI-agent frameworks with standard mechatronics for conveyor belt block-sorting tasks, combining an SO101 arm for color-based block recognition with two X-Arms utilizing an ultrasonics sensor.

The system hardware was modeled in Onshape® and SOLIDWORKS® and evaluated using physical experiments and computational simulations. For the X-Arm system, standard mechatronics programming applied forward and inverse kinematics alongside ultrasonic feedback for automated object detection and pick-and-place routines. For the SO101 arm, YOLO and OpenCV executed object detection and spatial tracking, while the IKPy library calculated kinematics to coordinate trajectories across the conveyor belt.

Experimental results demonstrated that standard kinematics performed effectively for highly repetitive picking, placing, and stacking 1-inch blocks. However, these methods relied heavily on initial calibration of the system, considering coordinates of the blocks as well as camera position. As a result, their applicability was largely limited to constrained and predictable environments. Conversely, AI-based approaches, including simulation training such as NVIDIA Isaac Sim and Isaac Lab, as well as simpler computer vision methods combining YOLO-based object detection with inverse kinematics, provided significantly greater flexibility. The environment was also connected to the SO101 arm through the ROS2 joint-state publisher, mirroring its movements in real time. These approaches enabled the system to identify and manipulate blocks without time-consuming calibration. Combining traditional kinematic control with AI-based methods such as Action Chunking with Transformers (ACT), Vision-Language models (VLMs), and Vision-Language-Action (VLAs) models yielded a more robust and adaptable robotic sorting framework than relying on conventional control alone.



46

Robotics & Embodied Intelligence



Liuyang Bai
Computer Science, B.S.
Spring 2027
NYU Shanghai

Latent World Model for Image Navigation

Image-goal navigation requires an embodied agent to reach a target location specified only by a goal image. Although recent navigation world models can predict future observations, planning directly in RGB space is computationally expensive and may limit scalability. This work investigates whether LeWorldModel (LeWM), a latent world model based on joint-embedding prediction, can provide useful and efficient planning signals for image-goal navigation.

We develop an experimental LeWM baseline and evaluate it on RECON and a subset of the Gibson dataset. The model is trained with a history size of four and one-step latent prediction, and is evaluated using trajectory candidates with a context size of four and a planning horizon of eight. We assess trajectory prediction accuracy, inference efficiency, and closed-loop navigation performance with cross-entropy-method planning.

On RECON, LeWM achieves an absolute trajectory error of 1.36 and a relative pose error of 0.34, demonstrating competitive trajectory prediction performance. On a single NVIDIA L40S GPU, evaluating 120 candidate trajectories requires 24.0 ms, corresponding to 0.20 ms per candidate and approximately 5,000 candidates per second. On held-out Gibson scenes, navigation from the source position achieves a success rate of 0.16 and an SPL of 0.13, while expert initialization reaches 0.82 for both metrics. Adding noise recovery and frontier exploration improves performance to a success rate of 0.26 and an SPL of 0.19.

These results show that latent world models are a feasible and computationally efficient baseline for image navigation. However, performance remains limited when the target is not visible, suggesting that better long-horizon training, distance-aware objectives, and exploration mechanisms are necessary for robust navigation.

NYU TANDON SCHOOL OF ENGINEERING

Latent World Model for Image Navigation

Liuyang Bai,[1] Xinhao Liu,[2] Jing Zhang,[2] Chen Feng[2]

[1] New York University Shanghai [2] New York University

Abstract

Image-goal navigation requires an agent to reach a location specified by a goal image. This project investigates whether LeWorldModel, or LeWM, can support navigation by predicting future states in a compact latent space rather than generating future RGB images [2]. We evaluate LeWM through trajectory prediction, inference-speed measurement, and closed-loop navigation. The results show that LeWM provides competitive prediction accuracy and fast candidate evaluation, but navigation remains difficult when the target is not visible.

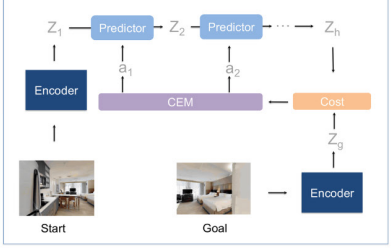
Background and Objective

Navigation World Models predict future observations conditioned on candidate actions and use these predictions to select a trajectory toward a goal [1]. More recent methods have explored dense visual representations and asymmetric world-action modeling for visual navigation [3, 4]. However, evaluating many trajectories in image space can be slow and difficult to scale. LeWM instead predicts future representations in a learned latent space [2]. The objective of this project is to determine whether this latent representation can support accurate trajectory prediction, efficient candidate evaluation, and closed-loop image-goal navigation. The current study is intended as an experimental baseline rather than a complete navigation system.

Method

LeWM encodes visual observations and predicts future latent representations conditioned on candidate actions [2]. During planning, candidate action sequences are sampled, predicted in latent space, and compared with the goal representation. The Cross-Entropy Method then selects the next action from the lowest-cost trajectories.

Figure 1: The latent-space planning pipeline adapted from the LeWorldModel architecture described by Maes et al. [2].



Results

On RECON, LeWM achieves an absolute trajectory error of 1.36 and a relative pose error of 0.34. On a single L40S GPU, the model evaluates 120 candidate trajectories in 24.00 milliseconds, which corresponds to approximately 4,999 candidates per second.

On unseen Gibson scenes, navigation from the original source position achieves a success rate of 0.16 and an SPL of 0.13. Expert initialization increases both metrics to 0.82. Adding noise recovery and frontier exploration improves source-position performance to a success rate of 0.26 and an SPL of 0.19.

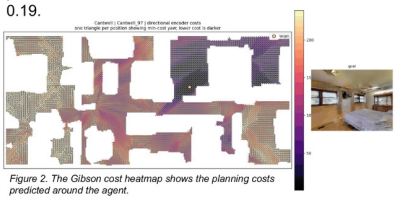


Figure 2. The Gibson cost heatmap shows the planning costs predicted around the agent.

Conclusion

LeWM is a feasible and computationally efficient baseline for local image-goal navigation. However, its performance decreases when the target is outside the current visual context. Future work should improve long-horizon prediction, distance-aware learning objectives, and exploration.

References

[1] A. Bar et al., "Navigation World Models," arXiv:2412.03572, 2025.
[2] L. Maes et al., "LeWorldModel: Stable End-to-End Joint-Embedding Predictive Architecture from Pixels," arXiv:2603.19312, 2026.
[3] M. Zhang et al., "RAE-NWM: Navigation World Model in Dense Visual Representation Space," arXiv:2603.09241, 2026.
[4] N. Yang et al., "WAM-Nav: Asymmetric Latent World-Action Modeling for Unified Visual Navigation," arXiv:2606.04907, 2026.

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.

Robotics & Embodied Intelligence



Ruoshui Li
Mechanical Engineering, B.S.
Spring 2028
NYU Tandon

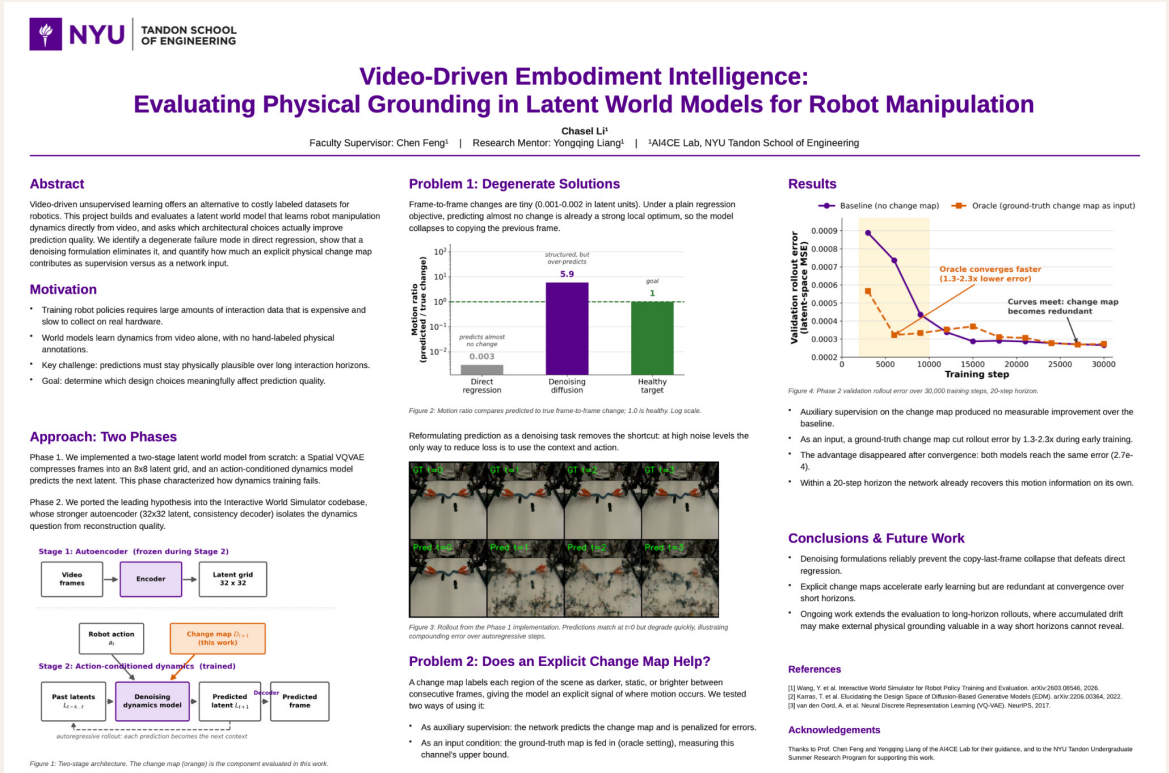
Video-Driven Embodiment Intelligence: A Progress Report

Introduction. Video-driven unsupervised learning offers an alternative to costly labeled datasets for robotics, by leveraging the structure already present in raw visual interaction data. Within this broader goal, this project focuses on a concrete instantiation: building and evaluating a latent world model that learns robot manipulation dynamics directly from video, without relying on hand-labeled physical annotations. The objective is to identify which architectural choices allow such a model to predict physically plausible future states over long interaction horizons.

Methods. Following the Interactive World Simulator framework, we train a two-stage model: a VQVAE compresses RGB observations into a compact latent space (Stage 1), and an action-conditioned dynamics model predicts future latents from past context and robot actions (Stage 2). We evaluated multiple dynamics formulations, including direct regression and a denoising diffusion approach based on EDM (Elucidating the Design Space of Diffusion-Based Generative Models), and tested whether an auxiliary "physical change map" between frames, capturing where the scene brightens or darkens, improves prediction, both as a training signal and as an explicit input to the network.

Results. Direct regression collapsed to a degenerate "copy last frame" solution, since true inter-frame changes are small in magnitude. Reformulating prediction as a denoising task eliminated this collapse and produced structured, action-responsive predictions. Supervising the change map as an auxiliary loss gave no measurable improvement, but providing the true change map as an input accelerated early training, an advantage that vanished over short (20-step) prediction horizons.

Conclusion. These findings clarify which design choices meaningfully affect video-driven dynamics learning and establish a reproducible testbed for evaluating physical grounding in world models. This research is ongoing; future work will extend this evaluation to long-horizon rollouts to determine whether explicit physical grounding provides lasting benefit for stable, label-free robot world models.



Robotics & Embodied Intelligence



Sabrina Addo
Computer Science, B.S.
Spring 2029
Spelman College

Understanding Communication Overhead and Accuracy Tradeoffs in Federated Learning Systems

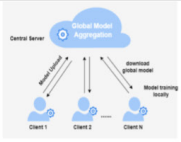
Federated learning enables machine learning models to be trained without centralizing users' data, making it promising for applications where privacy is a concern. However, its adoption remains limited, in part because of network communication overhead and uncertainty about task performance. The factors that shape the tradeoff between communication cost and model accuracy are not yet well understood or consistently documented. Our research reviews and organizes existing literature on key factors that may influence this tradeoff, including model size, data distribution, and client participation rates, and then conducts systematic experiments to evaluate their effects. By clarifying how these factors affect both efficiency and accuracy, this work aims to support better-informed design decisions and broader deployment of federated learning systems in privacy-sensitive domains.

Understanding Communication Overhead and Task Performance Tradeoffs in Federated

Sabrina Addo, Spelman College
Advised by Prof.Fraida Fund & Prof. Shivendra Panwar, NYU Tandon School of Engineering,

Introduction

Federated learning enables machine learning models to be trained **on users' own data while keeping it private.**

Federated Learning	Centralized Learning
	
Training on Device	Training on Cloud
Private Data	Centralized Data
Communication Overhead per Training Round	Minimal Overhead

Research gap: The factors that shape the tradeoff between communication cost and model accuracy are not yet well understood or consistently documented.

Objective

Review and **organize** existing literature on key factors that may influence communication/accuracy tradeoff.

Reproduce existing results.

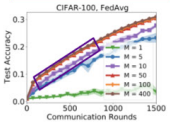
Evaluate using a new "communication overhead to target accuracy" metric.

Results

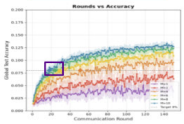
Client participation rates

Previous finding [1]: Increasing cohort size (M) improves convergence, but the gains diminish after a threshold.

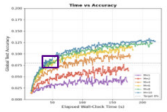
Their result: No change from **M=50→100→400**



Similarly in ours: no change from **M=6→8→10**



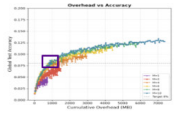
Time-to-accuracy gains worsen after M=6; Curves for **M=6,8,10** overlap significantly



Sharp drop in time from **M=1→6; Minimal additional speed** gains beyond M=6



Data transfer grows rapidly past **M=6 for minor accuracy gains**; Overhead cost outweighs convergence benefits at higher M



Total MB gradually drops from M=1→6 and plateaus beyond M=6; larger cohorts increase network load **without proportional accuracy benefits.**



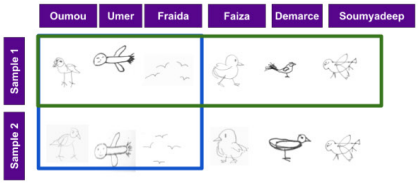
Conclusion

Paper Validated at Smaller Scale: Successfully reproduced Charles et al., proving accuracy gains do not scale linearly with cohort size and plateau around M=6.
Overhead Trade-Off Limit: Adding clients past M=6 yields the same accuracy in the same time, while continuing to add unnecessary communication overhead.

Classroom Materials

We adapt our research for classroom use, to demonstrate to undergraduate machine learning students how federated learning could benefit them personally.

Each student draws samples of 10 specified images (bird shown):



Students see that a model **trained on their own data** does better at their own future samples, than one **trained only on other users' data**. Federated learning supports this and privacy.

Future Work

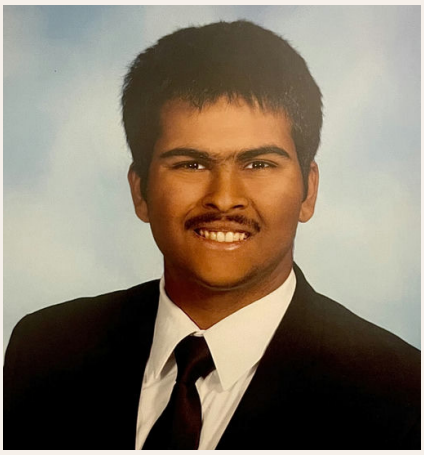
Evaluating the overhead–accuracy trade-off under varying factors:

- Model size
- Frequency of communication
- Data distribution (IID vs non-IID)
- Decentralized federated learning systems
- Device heterogeneity
- Communication scheduling [2]

References

[1] Charles, Zachary, et al. "On large-cohort training for federated learning." *Advances in neural information processing systems*34 (2021): 20461-20475.
[2] Nariman, G.S., Hamarashid, H.K. Communication overhead reduction in federated learning: a review. *Int J Data Sci Anal* **19**, 185–216 (2025)

Robotics & Embodied Intelligence



Tejas Siddaramaiah
Computer Science, B.S.
Fall 2027
NYU Tandon

Spatio-Temporal Mapping and LLM Timeline Analysis for FDNY Battery Fire Safety

Lithium-ion battery technology, electric vehicles (EVs), and micro-mobility devices pose emerging operational risks to modern urban fire safety. Continuing the igNYte laboratory's research framework established last year which evaluated historical general structural fire patterns across New York City, this project focuses specifically on battery-related fire incidents recorded by the Fire Department of the City of New York (FDNY). The objective is to evaluate battery fire dynamics, map spatial vulnerability, and benchmark incident progressions against general fires.

The study merges New York Fire Incident Reporting System (NYFIRS) records with daily weather datasets and operational metrics (alarm levels, units on scene, and suppression systems) to analyze fire dynamics and distinguish routine baseline events from rare, multi-alarm fires. Utilizing data science and GIS-based spatio-temporal mapping, the research identifies urban risk clusters, paying attention to building types and socioeconomically vulnerable populations. Additionally, a Large Language Model (LLM) pipeline is being developed to extract temporal events and compare battery fire timelines against general fire progressions.

Spatio-temporal mapping pinpoints critical risk clusters across New York City neighborhoods, detailing where battery fires concentrate relative to building characteristics and vulnerable populations. Operational analyses isolate the ten-codes within battery fires compared to standard incidents, while the LLM pipeline reconstructs incident progression timelines to highlight unique tactical challenges during response.

Primary outcomes include spatial risk maps to prioritize community education alongside empirical evidence to inform FDNY tactical response, specialized firefighter training, and updated building safety codes. Findings will be presented directly to the FDNY to improve fire safety.

NYU TANDON SCHOOL OF ENGINEERING

igNYte

Spatio-temporal Analysis of Battery Fire Data in New York City

Tejas Siddaramaiah,[1] Augustin Guibaud[2]

1. Research Assistant and CSE 2. Principal Investigator and MAE

Abstract

This project analyzes FDNY battery fire records, weather data, and operational metrics using GIS mapping. By identifying risk clusters and comparing battery fire dynamics to general fires, it delivers spatial risk maps and empirical insights to guide FDNY response, firefighter training, and updated building safety codes.

Background

Lithium-ion battery technology, electric vehicles (EVs), and micro-mobility devices pose emerging, high-intensity operational risks to modern urban fire safety. Continuing the igNYte laboratory's research framework established last year, which evaluated historical general fire patterns across New York City, this study focuses specifically on battery-related fire emergencies recorded in the Bureau of Fire Investigation (BFI) and New York Fire Incident Reporting System (NYFIRS) databases. The primary objective is to evaluate battery fire dynamics, map spatial vulnerability, and benchmark incident classifications against general fires to strengthen municipal emergency response.

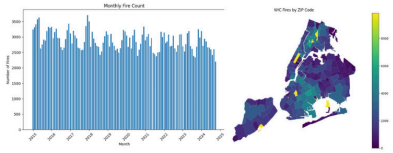


Figure 1: Previous spatio-temporal analysis on general New York City fires [1]

Methodology

The study analyzed 976 BFI records spanning July 2021 to August 2024. Merging BFI records with the NYFIRS incident database yielded 391 matched battery fires with complete operational response indicators (alarm levels, units on scene, floor elevation, etc.)

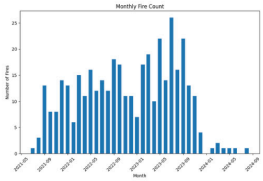


Figure 2: Monthly fire counts of lithium-ion battery fires

alongside daily weather metrics (temperature, humidity, precipitation, and wind speed). Utilizing data science techniques and GIS-based spatio-temporal mapping [2], the research evaluates weather independence and identifies urban risk clusters, paying dedicated attention to building types, occupancy, and socioeconomically vulnerable populations.

Results

Spatio-temporal mapping pinpoints critical risk clusters across New York City neighborhoods, detailing where battery fires concentrate relative to building characteristics and demographic vulnerabilities.

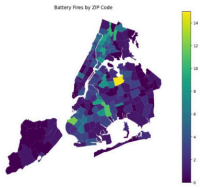


Figure 3: Spatial analysis of lithium-ion battery fires in New York City

Statistical modeling shows that daily battery fire occurrence operates largely independent of weather variations ($R^2 < 0.01$), and temporal analysis shows that mornings have a greater number of incidents. Spatial analyses isolate building types and ZIP codes, finding that Brooklyn and the Bronx contain most battery fires while neighborhoods with walk-up apartments and asylums and homes contain a greater number of battery fires.



Figure 4: Building class analysis with respect to ZIP code and lithium-ion battery fire counts

Conclusions

The study shows that daily battery fire occurrences operate largely independent of weather variations, unlike general fires. Battery fires are greater in Brooklyn and the Bronx and neighborhoods with walk-up apartments and asylums and homes, similar to general fires.

Future Work

We will continue our analysis of battery fires by creating a Large Language Model (LLM) pipeline to better understand the timelines of general and battery fires.

Works Cited

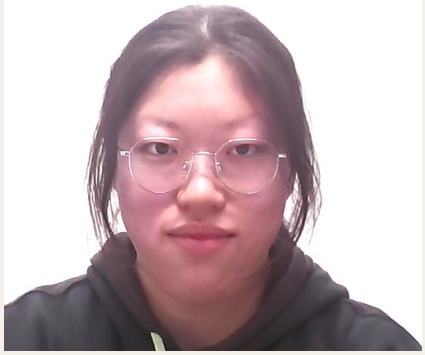
[1] Siddaramaiah, T., & Guibaud, A. (2025). Examining emerging fire risk in New York City using the FDNY incident database. NYU Tandon Undergraduate Summer Research Program Booklet, 44.

[2] Bonestroo, H., & Guibaud, A. (2026). Structural fires and emergency response in New York City: A spatiotemporal analysis of fire incidence and severity. Fire Safety Journal, 104718.

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project. Thanks to Hannah Bonestroo for pre-processed data and the FDNY for their continued knowledge and support.

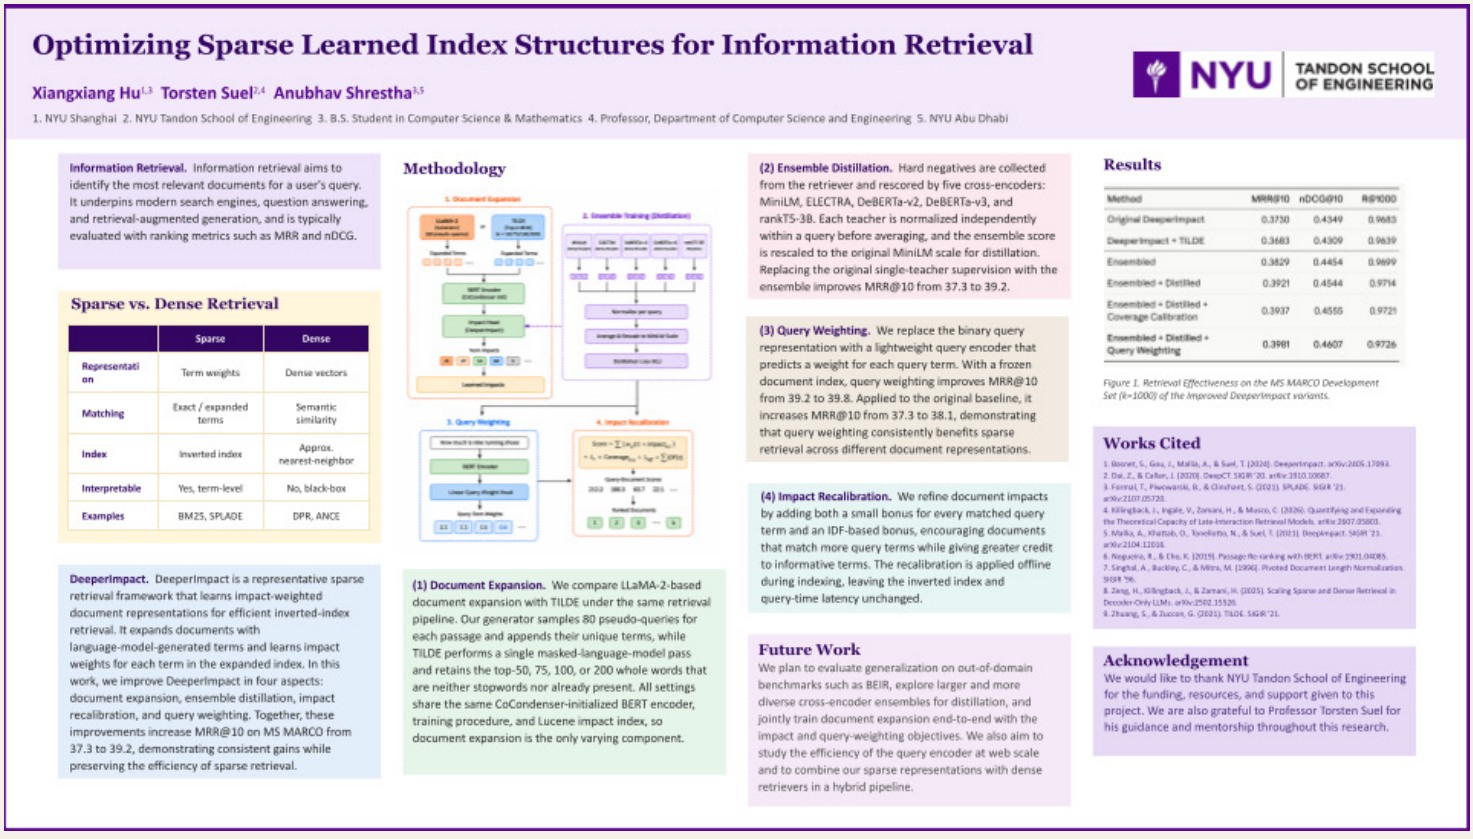
Robotics & Embodied Intelligence



Xiangxiang Hu
Computer Science, B.S.
Spring 2028
NYU Shanghai

Improving Sparse Retrieval: Expansion, Distillation, and Impact Calibration in DeeperImpact

Information retrieval seeks to identify and return the most relevant passages or documents in response to a user's query. Among neural approaches, sparse models represent text as high-dimensional sparse vectors whose nonzero entries are term-specific relevance scores, which keeps them compatible with the inverted-index infrastructure that already powers production search engines. DeeperImpact is one such framework. It first expands each document with terms that a language model predicts a user might search for, then learns an impact weight for every term in the expanded document. In this work, we improve DeeperImpact along four directions. First, we study document expansion, comparing LLaMA-based query generation against the far cheaper TILDE lexical expansion under a matched training recipe. Second, we strengthen supervision by distilling from an ensemble of five cross-encoder teachers together with self-distillation on mined hard negatives. Third, we recalibrate the learned impact weights after training. The additive scoring function systematically over-ranks documents that match only part of a query, and a single coverage parameter baked directly into the index corrects this at no cost during query processing. Fourth, we move query term weighting into the model itself, jointly training a query-side weighting component with the document-side impact model so that both sides are optimized end to end under the same ranking objective, rather than relying on a fixed heuristic such as IDF. Together, these improvements raise MRR@10 on MS MARCO from 37.3 to 39.2.



Robotics & Embodied Intelligence



Xingtong Liu
Mathematics, B.S.
Spring 2028
NYU Shanghai

Learning-based Robotic Rope Manipulation

Robotic manipulation of deformable linear objects (DLOs), such as ropes, remains a challenging problem due to their complex physical behavior and evolving topological configurations. Reliable rope manipulation is critical for applications including cable routing, surgical assistance, industrial automation, and household robotics. This project explores current approaches to rope representation, physics-based simulation, and learning-based manipulation to understand these challenges better. A systematic literature review was conducted on rope perception, state representation, dynamics modeling, and reinforcement learning for robotic manipulation. NVIDIA Newton was also explored as a physics-based simulation platform to investigate rope modeling, physical parameterization, and simulation environment design. The study summarizes the strengths and limitations of existing approaches and identifies several open research challenges, including representing rope topology, modeling deformable dynamics, and integrating physics simulation with learning-based manipulation. The simulation environment established during this project provides a foundation for future implementation and evaluation of robotic rope manipulation methods. This work establishes a foundation for future research on learning-based rope manipulation and contributes toward the development of more reliable robotic systems capable of manipulating deformable objects in complex real-world environments.

NYU TANDON SCHOOL OF ENGINEERING

Learning-Based Robotic Rope Manipulation

Simulation, Representation, and Control for Knot-Tying Tasks

Xingtong Liu | NYU Tandon School of Engineering | Advisor: Yongqing Liang | Professor: Chen Feng

Abstract

Robotic manipulation of deformable linear objects (DLOs), such as ropes and cables, remains challenging because their shape changes continuously under contact, friction, and self-collision. Knot tying further requires reasoning about both geometry and topology, including over-under crossing relations.

This project studies learning-based rope manipulation through literature review, physics-based simulation, and representation analysis. We reviewed Cosserat rod modeling, graph-based neural dynamics, object-centric sim-to-real RL, diffusion policy, MPC, and reinforcement learning. In parallel, NVIDIA Newton was explored as a simulation platform for rope modeling, parameter tuning, offline rendering, and knot initialization.

The current goal is to build a simulation foundation for future learning-based knot manipulation policies that can generalize across rope materials, initial states, and knot types.

Research Question

How can physics-based simulation and learning-based control be combined to enable robust robotic rope manipulation and knot tying?

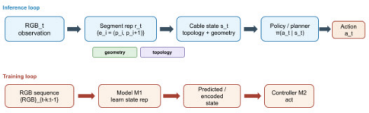
Motivation

- **Deformation:** rope states are high-dimensional and history-dependent.
- **Contact:** friction, self-collision, and grasping change the dynamics.
- **Topology:** knots depend on crossing order, not only point positions.
- **Generalization:** policies should transfer across materials, rope lengths, and tasks.

Previous Work

Category	Examples	Takeaway / Relevance
Perception / state estimation	TrackDLO GraphDLO	Rope state can be represented with tracked points, graphs, or estimated geometry.
Learned dynamics / world models	AdaptiGraph Particle-Grid Dynamics Particle Former Baseline: GBND	Compare material adaptation, particle-based dynamics, and scene-specific assumptions.
Planning / routing policies	BaB-ND Physics-informed planning Goal-conditioned simplification Cable routing with RL / Diffusion	Long-horizon manipulation needs planning over contacts, goals, and constraints.
Knot tying / topology	RoboHitch TWISTED-RL / TWISTED DexKnot Bi-manual shoe lacing	Knot tasks require reasoning about crossings, topology, and multi-step manipulation.

Methods



Cable State Representation

RGB observation

Centerline points

Ordered points p_i

Segments e_i=(p_i,p_i+1)

Cable state s_t

We intended to use segments to represent the rope and use GIF demonstrations from the Animated Knots (3) website as our data source.

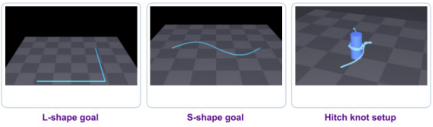


Double Fisherman's bend knot: https://www.animatedknots.com/double-fishermans-bend-knot

Fig8: https://www.animatedknots.com/figure-8-knot

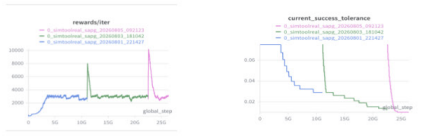
Bowline: https://www.animatedknots.com/bowline-knot

Newton Experiments



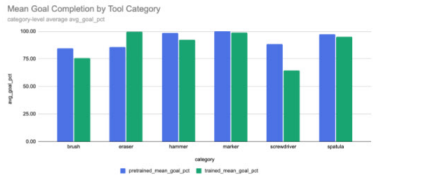
L-shape goal | S-shape goal | Hitch knot setup

SimToolReal Reproduction



Reward stabilizes after early learning and curriculum transitions.

Tolerance tightens over training, so later success requires higher precision.



Future Work

Build stable Newton environments for basic knot manipulation.
Refine RGB centerline extraction and segment-based cable state representation.
Tune friction, damping, stiffness, and segment length across materials.
Generate rollouts for object-centric RL, MPC, Diffusion Policy, or PPO training.
Evaluate generalization across rope materials, initial states, and knot types.

References

[1] Kedia, K., Lum, T. G. W., Bohg, J., & Liu, C. K. (2026). *SimToolReal: An object-centric policy for zero-shot dexterous tool manipulation*. arXiv. https://arxiv.org/abs/2602.16863
[2] Newton Project. (2026). *Newton: An open-source, GPU-accelerated physics simulation engine* [Computer software]. GitHub. https://github.com/newton-physics/newton
[3] Grog LLC. (n.d.). *Animated Knots by Grog*. https://www.animatedknots.com/

Acknowledgement: NYU Tandon School of Engineering; Advisor: Yongqing Liang; Professor: Chen Feng.

Robotics & Embodied Intelligence



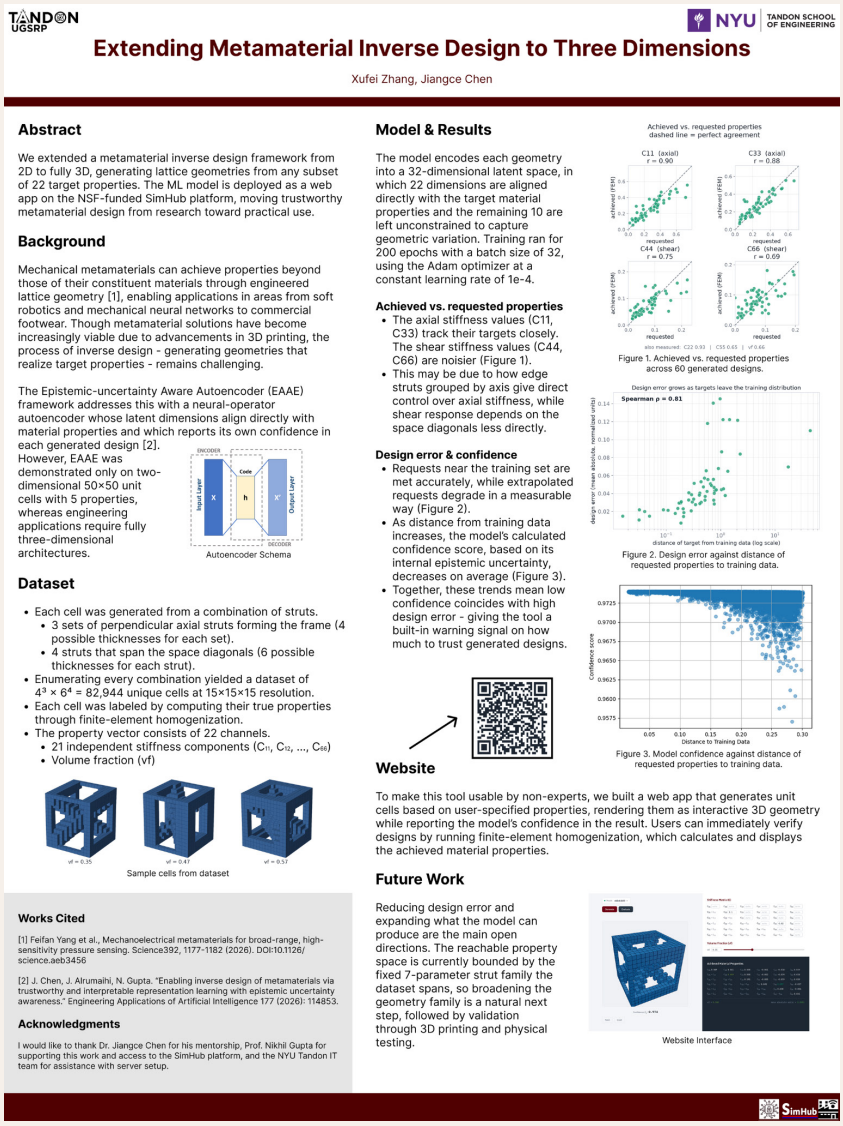
Xufei Zhang
Computer Science, B.S.
Spring 2028
NYU Tandon

Extending Trustworthy Metamaterial Inverse Design to Three Dimensions: An Accessible, Uncertainty-Aware Web Tool

Mechanical metamaterials can achieve properties beyond those of their constituent materials through engineered lattice geometry, enabling applications in areas from soft robotics and mechanical neural networks to commercial footwear. Though metamaterial solutions have become increasingly viable due to advancements in 3D printing, the process of inverse design - generating geometries that realize target properties - remains challenging. The Epistemic-uncertainty Aware Autoencoder (EAAE) addresses this with a neural-operator autoencoder whose latent dimensions align directly with material properties and which reports its own confidence in each generated design. However, EAAE was demonstrated only on two-dimensional 50×50 unit cells with 5 properties, whereas engineering applications require fully three-dimensional architectures.

This project extended the EAAE framework to 3D and made it usable by non-experts. Because the 3D elasticity tensor is a symmetric 6×6 matrix, the property vector grew from 5 to 22 channels (21 independent stiffness components + volume fraction). We generated a labeled dataset of 82,944 volumetric unit cells at 15×15×15 resolution, computing the true properties of each cell by finite -element homogenization distributed across an HPC cluster, and trained the model on it. Geometries are produced by optimizing the unconstrained latent dimensions to maximize the model's familiarity with the requested target properties. The resulting model generates connected 3D lattice geometries from any user-specified subset of the 22 properties, automatically filling the remainder, and reports a confidence score that flags out-of-distribution requests. We built a web interface, backed by a containerized inference server, that lets users specify target properties, generate a geometry, inspect it in 3D, and verify the achieved properties through on-demand finite-element homogenization.

By extending a trustworthy inverse-design framework to three dimensions and delivering it as an accessible tool on the NSF-funded SimHub platform, this work moves uncertainty-aware metamaterial design from academic research toward practical use.



TANDON UGSRP

NYU TANDON SCHOOL OF ENGINEERING

Extending Metamaterial Inverse Design to Three Dimensions

Xufei Zhang, Jiangce Chen

Abstract

We extended a metamaterial inverse design framework from 2D to fully 3D, generating lattice geometries from any subset of 22 target properties. The ML model is deployed as a web app on the NSF-funded SimHub platform, moving trustworthy metamaterial design from research toward practical use.

Background

Mechanical metamaterials can achieve properties beyond those of their constituent materials through engineered lattice geometry [1], enabling applications in areas from soft robotics and mechanical neural networks to commercial footwear. Though metamaterial solutions have become increasingly viable due to advancements in 3D printing, the process of inverse design - generating geometries that realize target properties - remains challenging.

The Epistemic-uncertainty Aware Autoencoder (EAAE) framework addresses this with a neural-operator autoencoder whose latent dimensions align directly with material properties and which reports its own confidence in each generated design [2]. However, EAAE was demonstrated only on two-dimensional 50×50 unit cells with 5 properties, whereas engineering applications require fully three-dimensional architectures.

Autoencoder Schema

Dataset

- Each cell was generated from a combination of struts.
  - 3 sets of perpendicular axial struts forming the frame (4 possible thicknesses for each set).
  - 4 struts that span the space diagonals (6 possible thicknesses for each strut).
- Enumerating every combination yielded a dataset of $4^3 \times 6^4 = 82{,}944$ unique cells at 15×15×15 resolution.
- Each cell was labeled by computing their true properties through finite-element homogenization.
- The property vector consists of 22 channels.
  - 21 independent stiffness components (C_{11}, C_{12}, ..., C_{66})
  - Volume fraction (vf)

Sample cells from dataset

Works Cited

[1] Feifan Yang et al., Mechanoelectrical metamaterials for broad-range, high-sensitivity pressure sensing. Science392, 1177-1182 (2026). DOI:10.1126/science.aeb3456

[2] J. Chen, J. Alrumaihi, N. Gupta. "Enabling inverse design of metamaterials via trustworthy and interpretable representation learning with epistemic uncertainty awareness." Engineering Applications of Artificial Intelligence 177 (2026): 114853.

Acknowledgments

I would like to thank Dr. Jiangce Chen for his mentorship, Prof. Nikhil Gupta for supporting this work and access to the SimHub platform, and the NYU Tandon IT team for assistance with server setup.

Model & Results

The model encodes each geometry into a 32-dimensional latent space, in which 22 dimensions are aligned directly with the target material properties and the remaining 10 are left unconstrained to capture geometric variation. Training ran for 200 epochs with a batch size of 32, using the Adam optimizer at a constant learning rate of 1e-4.

Achieved vs. requested properties

- The axial stiffness values (C11, C33) track their targets closely. The shear stiffness values (C44, C66) are noisier (Figure 1).
- This may be due to how edge struts grouped by axis give direct control over axial stiffness, while shear response depends on the space diagonals less directly.

Figure 1. Achieved vs. requested properties across 60 generated designs.

Design error & confidence

- Requests near the training set are met accurately, while extrapolated requests degrade in a measurable way (Figure 2).
- As distance from training data increases, the model's calculated confidence score, based on its internal epistemic uncertainty, decreases on average (Figure 3).
- Together, these trends mean low confidence coincides with high design error - giving the tool a built-in warning signal on how much to trust generated designs.

Figure 2. Design error against distance of requested properties to training data.

Figure 3. Model confidence against distance of requested properties to training data.

Website

To make this tool usable by non-experts, we built a web app that generates unit cells based on user-specified properties, rendering them as interactive 3D geometry while reporting the model's confidence in the result. Users can immediately verify designs by running finite-element homogenization, which calculates and displays the achieved material properties.

Future Work

Reducing design error and expanding what the model can produce are the main open directions. The reachable property space is currently bounded by the fixed 7-parameter strut family the dataset spans, so broadening the geometry family is a natural next step, followed by validation through 3D printing and physical testing.

Website Interface



Engineering Health

Engineering and Health focuses on seamless integration of engineering with biology, clinical practice, data science, and AI to create innovative solutions for maintaining wellness, preventing disease, and treating illness more effectively, economically, and equitably.

NYU is forging a transformative path in health engineering by converging its exceptional strengths in engineering, medicine, basic sciences, math, data science, and dentistry. Spearheaded by the cross-disciplinary Institute for Engineering Health — co-led by Tandon and NYU Langone with a dynamic dual presence in Brooklyn and Manhattan along with strategic global partnerships — NYU is tackling critical challenges in immunoengineering and biomedical engineering with a focus on translation, scalability, and societal impact.

Engineering Health



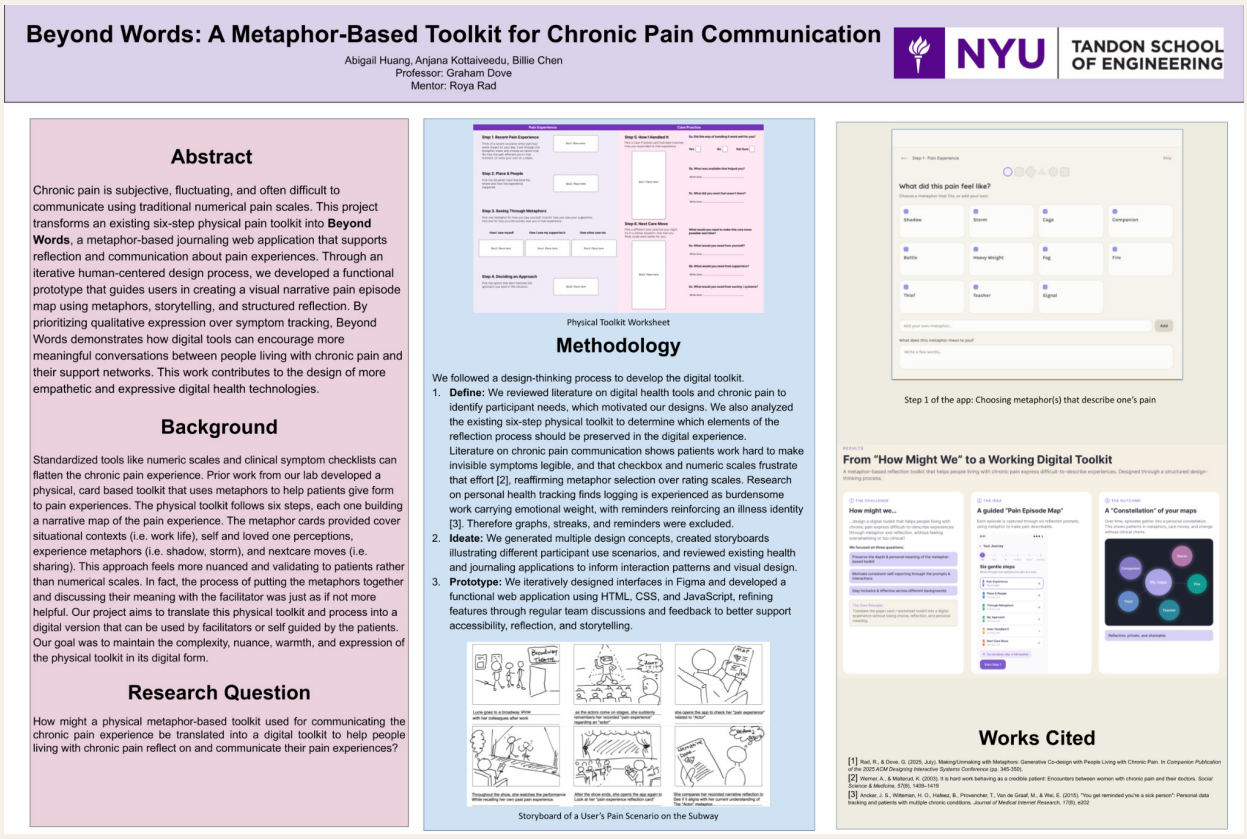
Abigail Huang
Computer Science, B.S.
Spring 2027
Stony Brook University



Anjana Kottaiveedu
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

Beyond Words: Translating a Metaphor-Based Toolkit into a Narrative Pain Episode Map for Chronic Pain Communication

The chronic pain experience is subjective, often fluctuating and hard to communicate. The typical use of numerical scales and symptom trackers do not capture the complexity, emotional impact, or the personal understanding of one's pain. Prior work in our lab shows that metaphor-based toolkits help provide form to otherwise invisible aspects of pain. This project aims to transform an existing six-step physical pain toolkit into Beyond Words, a metaphor-based journaling web application that supports pain reflection and communication between participants and facilitators. We followed the design-thinking process, which consists of empathizing, ideating, and prototyping, to make decisions about the app. We used key principles to iteratively prototype application layouts in Figma. This was coupled with storyboards to understand the reasons why participants would interact with the toolkit, along with the takeaways. The resulting toolkit, Beyond Words, is a pain episode narrative map that guides users through six steps to produce a visual record of their experience. This is a functional prototype that preserves the expressive qualities of the original physical toolkit while improving accessibility, organization, and ease of use in a digital format. The map allows participants to select between session modes, elaborate upon selected metaphors, an archive with constellation views of their episode, and options to share. The design emphasizes self-reflection, storytelling, and communication rather than symptom quantification alone. The digital toolkit strays away from graphs, streaks, and reminders to emphasize the qualitative reflections. By supporting more expressions of pain experiences, Beyond Words demonstrates how digital tools can complement traditional pain assessment methods and foster more meaningful conversations between individuals with chronic pain and support networks. This work contributes to ongoing research on designing more empathetic and expressive digital health technologies.





Anika Das
Biomedical Engineering, B.S.
Winter 2027
University of Michigan
AHA Fellow

Modeling CAR T-Cell Therapy in Pancreatic Ductal Adenocarcinoma using an Organ-on-a-Chip Platform

Pancreatic ductal adenocarcinoma (PDAC) is the most common form of pancreatic cancer accounting for 95% of all cases [1]. Due to the aggressive nature of PDAC and difficulty diagnosing at an early stage, 80% of cases are identified at an advanced stage leading to a 5
year survivorship rate of less than 13%[2]. CAR T-cell therapy is a highly promising treatment which reprograms a patient's own T-cells to recognize and target cancer cells by attaching a gene for a Chimeric Antigen Receptor (CAR) to T-cells [3]. While CAR T-cell therapies have been highly effective in treating blood cancers such as leukemia, resulting in high remission rates, applications to solid tumors such as in PDAC had proven difficult as CAR T-cell infiltration is difficult due to the large stroma [4]. In order to recapitulate the tumour microenvironment in PDAC, our lab utilized organ-on-a-chip platforms to co-culture patient derived cancer cell organoids (PDOs) and stromal cells with cancer associated fibroblasts (CAFs). We then added CAR T-cells and measured changes in tumor organoid diameter using an imaging processing pipeline. Our results validated the efficacy of CART therapy on PDAC solid tumors and demonstrated the infiltration of CAR T-cells in the stroma. Future directions include utilizing these organ-on-a-chip platforms to test combination therapies and treatment resistance on patient derived samples as well as further developing our AI model to predict treatment responsiveness and improve upon patient outcomes.

NYU TANDON SCHOOL OF ENGINEERING

BIOMEDICAL ENGINEERING UNIVERSITY OF MICHIGAN

Modeling CAR T-Cell Therapy in Pancreatic Ductal Adenocarcinoma using an Organ-on-a-Chip Platform

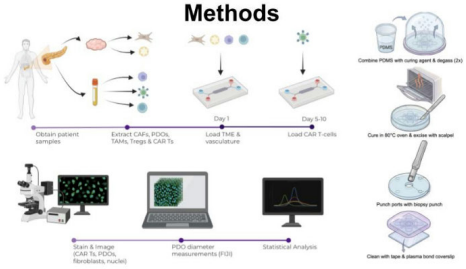
Anika Das,[1] David Baek,[2] Katsuo Kurabayashi [3]

1. Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI
2. Department of Biomedical Engineering, Tandon School of Engineering, New York University, Brooklyn, NY
3. Department of Mechanical and Aerospace Engineering, Tandon School of Engineering, New York University, Brooklyn, NY

Abstract

Pancreatic ductal adenocarcinoma (PDAC) is the most common form of pancreatic cancer accounting for 95% of all cases [1]. Due to the aggressive nature of PDAC and difficulty diagnosing at an early stage, 80% of cases are identified at an advanced stage leading to a 5 year survivorship rate of less than 13% [2]. CAR T-cell therapy is a highly promising treatment which reprograms a patient's own T-cells to recognize and target cancer cells by attaching a gene for a Chimeric Antigen Receptor (CAR) to T-cells [3]. While CAR T-cell therapies have been highly effective in treating blood cancers such as leukemia, resulting in high remission rates, applications to solid tumors such as in PDAC had proven difficult as CAR T-cell infiltration is difficult due to the large stroma [4]. In order to recapitulate the tumour microenvironment in PDAC, our lab utilized organ-on-a-chip platforms to co-culture patient derived cancer cell organoids (PDOs) and stromal cells with cancer associated fibroblasts (CAFs). We then added CAR T-cells and measured changes in tumor organoid diameter using an imaging processing pipeline. Our results validated the efficacy of CART therapy on PDAC solid tumors and demonstrated the infiltration of CAR T-cells in the stroma. Future directions include utilizing these organ-on-a-chip platforms to test combination therapies and treatment resistance on patient derived samples as well as further developing our AI model to predict treatment responsiveness and improve upon patient outcomes.

Methods



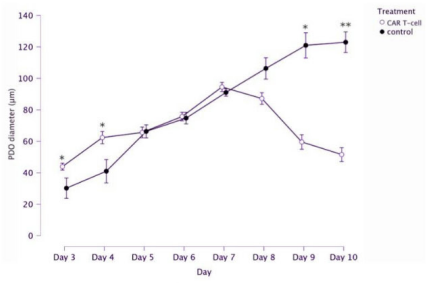
In order to fabricate the microfluidic devices, PDMS gel base and curing agent were mixed at a 10:1 ratio, degassed under a vacuum, poured into microfluidic molds, and degassed again to eliminate residual air bubbles. The molded PDMS was thermally cured at 80°C for ~1.5 hours, and excised using a scalpel. Outer reservoir ports (4 mm) and central microchannel inlets (1–1.4 mm) were created using biopsy punches. Following tape cleaning to remove particulate debris, devices and glass coverslips were exposed to oxygen plasma for 2.5 minutes, immediately bonded, and baked overnight at 80°C to finalize covalent bonding.

Patient derived tumor tissues and blood samples were processed to isolate cancer associated fibroblasts (CAFs), patient derived organoids (PDOs), immune cells (TAMs, Tregs), and engineered CAR T-cells. On Day 1, PDOs, vascular cells and immune cells were seeded into the central microchannels of the fabricated devices to mimic the tumor microenvironment (TME). CAFs were seeded along the inner and outer rings of the device and vascular cells were seeded between the rings to mimic the solid tumor stroma. After establishing organoid cultures, CAR T-cells were introduced on Day 5. Microfluidic platforms were maintained with twice daily media replacements and imaged daily. After day 10, devices underwent immunostaining to visualize CAR T-cell infiltration, PDO morphology, and cell nuclei. Organoid diameters were quantified across time points using ImageJ/FIJI processing software, and statistical significance across treatment groups was evaluated using ANOVA testing.

Results

To evaluate the therapeutic efficacy of CAR T-cells in our solid tumor microenvironment model, patient-derived organoid (PDO) diameters were monitored daily from Day 3 through Day 10. Both control and CART-treated PDOs exhibited steady baseline growth from Day 3 to Day 7, reaching similar maximum diameters (~90–95 µm). Following CAR T-cell introduction, treated PDOs demonstrated a marked divergence starting at Day 8, experiencing progressive shrinkage down to an average diameter of ~50 µm by Day 10. In contrast, untreated control organoids continued expanding, exceeding 120 µm by Day 10. These results demonstrate that CAR T-cells successfully infiltrate the dense microfluidic stroma to exert delayed, effective tumor killing and organoid regression by Days 9 and 10.

Figure 1: PDO Diameter as a Result of CAR T-cell Therapy



Discussion

This study validates a patient-derived microfluidic organ-on-a-chip platform for modeling CAR T-cell therapy in pancreatic ductal adenocarcinoma. The device successfully recapitulates the PDAC solid tumor stroma to evaluate immune cell infiltration and PDO response. The initial lag in organoid regression (Days 3–7) reflects the time required for CAR T-cells to infiltrate the stromal barrier, while significant size reduction by Day 10 confirms targeted killing after infiltration. Ultimately, this platform provides a controllable, high-throughput tool to bridge 2D models and animal studies for evaluating patient-specific treatment efficacy. Future studies plan to investigate the efficacy of CAR T-cell combination therapies on PDAC.

Sources

[1] Sarantis P. et al. *World J. Gastrointest. Oncol.* **12**(2), 173–181 (2020).
[2] Sesock C. *Pancreatic Cancer Action Network* (2025).
[3] Fonkoua L. A. K. et al. *Mol. Ther. Oncolytics* 25, 69–77 (2022).
[4] National Cancer Institute. *NCI Dictionary of Cancer Terms* (n.d.).

Acknowledgement

I would like to thank my mentor, David Baek and PI, Katsuo Kurabayashi for their support throughout this project.

American Heart Association.

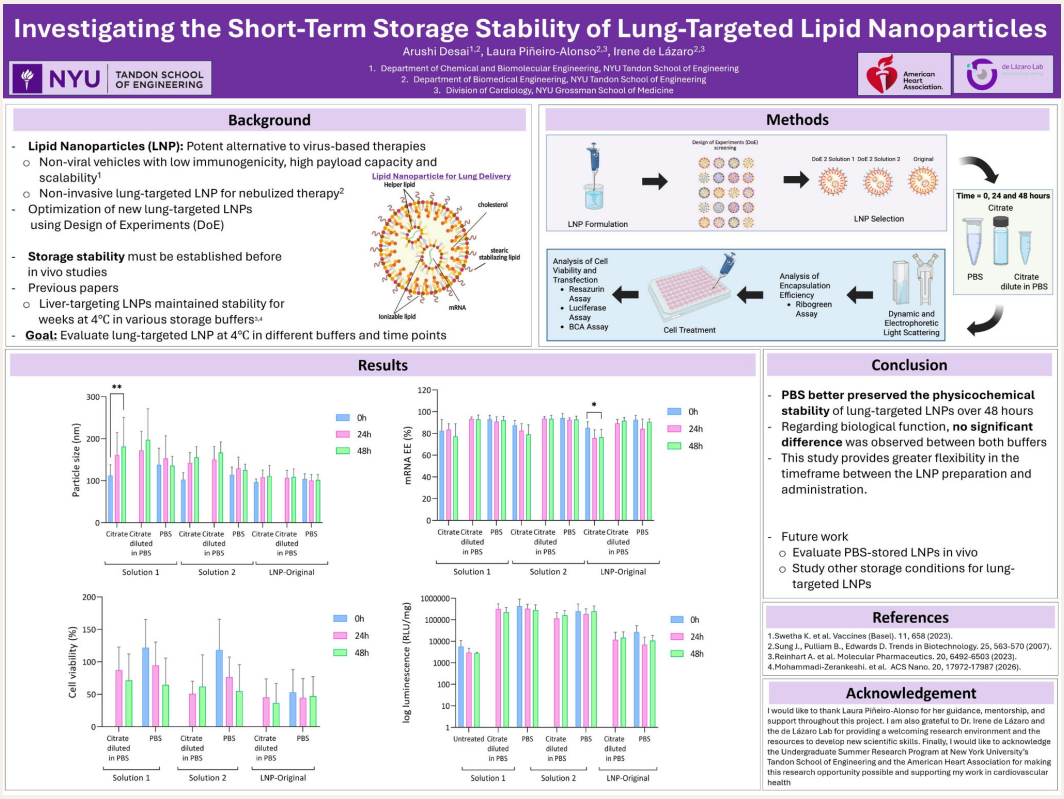
NYU TANDON UGSRP



Arushi Desai
Biomolecular Science, B.S.
Spring 2029
NYU Tandon
AHA Fellow

Investigating the Short-Term Storage Stability of Lung-Targeted Lipid Nanoparticles

Lipid Nanoparticles (LNP) are a promising non-viral delivery system for mRNA therapeutics due to their low immunogenicity, high payload capacity, scalability and compatibility with non-invasive pulmonary delivery. Accordingly, we developed lung-targeting LNPs as an alternative platform to adeno-associated virus (AAV)-mediated Fibroblast Growth Factor 21 for arrhythmogenic right ventricular cardiomyopathy (ARVC). Using a Design of Experiment (DOE) methodology, we have identified lung-targeting LNPs with optimal cell viability and transfection efficiency. To progress into in vivo studies, it is important to understand the storage stability of these LNPs. While previous studies have demonstrated that liver-targeting LNPs remain stable at 4 degrees Celsius in multiple storage buffers, the stability of these lung-targeting prototype LNPs has not been established. This study evaluates how storage buffer and time after formulation affect the physicochemical properties and transfection capability of our final prototypes. The lead DOE candidates were formulated with luciferase-encoding mRNA, stored at 4 degrees Celsius in PBS or citrate buffer, and administered to A549 cells at 0, 24 and 48 hours after formulation. Particle size and zeta potential were measured using dynamic and electrophoretic light scattering, respectively. The ribogreen assay was used to evaluate encapsulation efficiency and mRNA concentration before treating the cells at each time point. Cell viability and transfection efficiency were measured a day after treatment using resazurin assay, bicinchoninic acid assay and luciferase assay. Results indicate that PBS provided more consistent physicochemical stability over 48 hours, whereas citrate led to greater variability in particle size and encapsulation efficiency. Nonetheless, transfection and cell viability remained similar regardless of storage buffer and time, indicating that short-term changes in physicochemical properties did not compromise LNP function. These findings suggest that our lung-targeted prototypes can be reliably stored in PBS and used within 48 hours of formulation without affecting performance. These storage conditions increase experimental flexibility in conducting reproducible in vivo studies. Our results advance the development of lung-targeting LNPs as a therapeutic for ARVC.

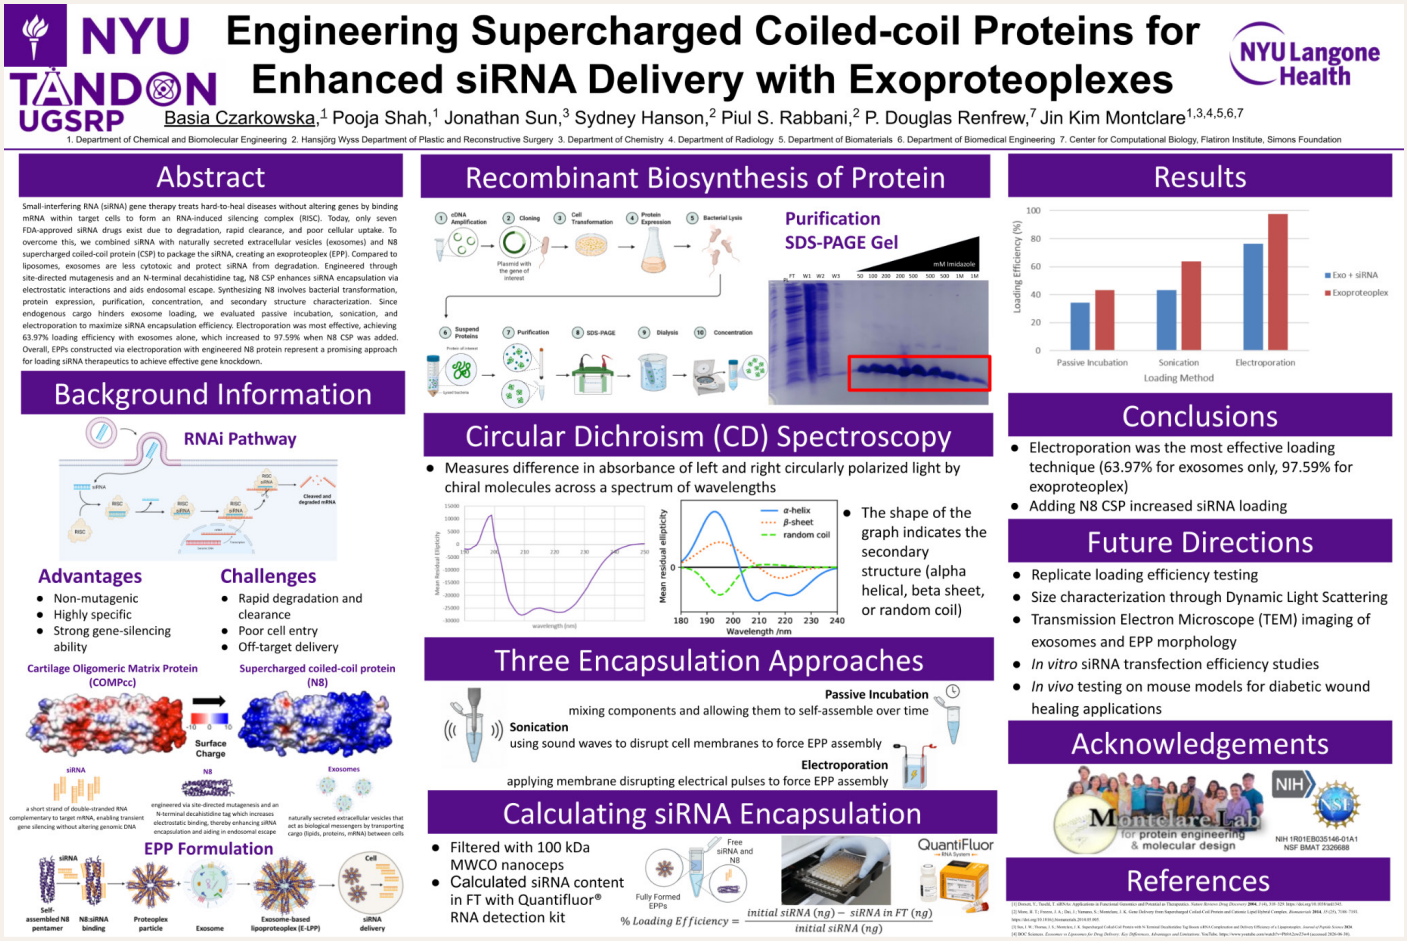




Barbara Czarkowska
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon
Thompson-Bartlett Fellow

Engineering Supercharged Coiled-coil Proteins for Enhanced siRNA Delivery with Exoproteoplexes

Small-interfering RNA (siRNA) gene therapy is a novel technology designed to treat patients with diseases that cannot be healed through conventional drugs. SiRNA is a short strand of double-stranded RNA complementary to target mRNA, enabling transient gene silencing without altering genomic DNA by incorporating into the RNA-induced silencing complex (RISC). Today, there are only seven FDA-approved siRNA drugs due to limitations such as degradation, rapid clearance, and poor cellular uptake. To combat this, we combined the siRNA with naturally secreted extracellular vesicles (exosomes) and N8 supercharged coiled-coil protein (CSP) to package the siRNA, creating an exoproteoplex (EPP). Compared to liposomes, exosomes are less cytotoxic delivery agents that protect siRNA from degradation, allowing it to enter the cell. We engineered the N8 protein via site-directed mutagenesis and an N-terminal decahistidine tag addition. These modifications increase electrostatic binding, thereby enhancing siRNA encapsulation, and aid in endosomal escape. Synthesizing this protein involves bacterial transformation, protein expression, purification, and concentration followed by secondary structure characterization. One challenge with exosome delivery vehicles is that they have their own endogenous cargo, making it difficult to load therapeutic siRNA. This issue, along with wanting to load more siRNA into the vehicle to increase therapeutic efficacy, are the reasons we tested encapsulation efficiency. Three loading methods were compared: passive incubation, sonication, and electroporation. Preliminarily, electroporation was found to be the most effective at encapsulating the siRNA with 63.97% loading efficiency for exosomes and siRNA alone, which increased to 97.59% after the addition of the N8 CSP. This study shows that EPPs constructed through electroporation with the engineered N8 protein are a promising way to load siRNA therapeutics for effective gene knockdown.





Cansu Altintas
Chemical and Biomolecular Engineering, B.S.
Spring 2027
NYU Tandon

Expanding the repertoire of solid cancer antigens using a novel inside-out logic-gated CAR T cell strategy

Chimeric antigen receptor (CAR) T-cell therapy has been notable in cancer treatments and is continuing to be a promising solution. It is highly successful in liquid tumors because it targets lineage-specific antigens, allowing both cancerous and healthy cells to be depleted. This strategy works with B cells as depletion of B cells is well tolerated, but it cannot be applied to other tissues as it can result in life-threatening conditions. Another reason why CAR-T therapy shows a limited effect on solid tumors is due to the lack of known tumor-specific antigens and the risk of on-target, off-tumor toxicity (OTOT). To solve this issue, the Yarmarkovich lab has pioneered a new class of CARs called peptide-centric (PC) CAR-T. They are able to recognize and target intracellular drivers as peptides presented on human leukocyte antigens (pHLA), which significantly expands the repertoire of targetable tumor antigens. Another strategy to increase the number of targetable antigens is the use of IF-gated logic CARs. IF-THEN-gated systems like SynNotch require recognition of two antigens to activate T cells. These systems avoid OTOT by targeting antigen pairs that are exclusively expressed on cancer cells and mutually excluded on healthy tissues. However, these strategies alone do not overcome the limited availability of tumor-specific targets. Cancer-specific intracellular antigens are not always available, and relying solely on logic-gated membrane antigens severely limits the pool of optimal antigen pairs. To overcome these limitations, we developed an inside-out strategy that combines logic-gated CARs with PC-CARs, enabling recognition of both intracellular and extracellular proteins while preventing OTOT . With this novel technology, we will be able to greatly expand the range of targetable antigens while keeping cancer, allowing development of new therapies for many solid cancers with unmet medical needs.

NYU TANDON SCHOOL OF ENGINEERING — NYU Langone Health

Expanding the repertoire of solid cancer antigens using a novel inside-out logic-gated CAR T cell strategy

Cansu Altintas,[1] Michele Palamenghi,[2] Mark Yarmarkovich[3]
1. Undergraduate Researcher at NYULH & NYU Tandon 2. Post Doctoral Fellow at the Yarmarkovich Lab 3. Principal Investigator at the Yarmarkovich Lab

Abstract

Chimeric antigen receptor (CAR) T-cell therapy has been notable in cancer treatments and is continuing to be a promising solution. However, CAR-T therapy shows a limited effect on solid tumors is due to the lack of known tumor-specific antigens and the risk of on-target, off-tumor toxicity (OTOT). To overcome these limitations, we developed an inside-out strategy that combines logic-gated CARs with peptide-centric (PC)-CARs. PC-CARS will be able to recognize and target intracellular drivers which will expand the available antigens to target. With an AND logic gate, OTOT will be prevented by increasing specificity by requiring recognition of two antigens for activation.

Background

CAR T-cell therapy is a form of immunotherapy in which a patient's own T cells are genetically engineered to target cancer cells. It is highly effective in treating liquid tumors because it targets lineage-specific antigens, allowing both cancerous and healthy cells to be depleted. This strategy works with B cells, as depletion of B cells is well tolerated. But it cannot be applied to other tissues as it can result in life-threatening conditions. Thus, lack of known tumor-specific antigens and the risk of OTOT are one of the biggest limitations in targeting solid tumors with CAR T-cells. The Yarmarkovich lab has pioneered a novel strategy to expand the number of targetable antigens, by developing a new class of CARs called peptide-centric (PC) CAR, which recognizes and targets intracellular drivers as peptides presented on human leukocyte antigens (pHLA) [1]. An anti-PHOX2B PC-CAR entered clinical trial in early 2026 to treat pediatric neuroblastoma patients. Another strategy to increase the number of targetable antigens is the use of AND-gated logic CARs. AND-gated systems like SynNotch require recognition of two antigens to activate T cells. These systems avoid OTOT by targeting antigen pairs that are exclusively co-expressed on cancer cells.

However, these strategies alone do not overcome the limited availability of tumor-specific targets. Cancer-specific intracellular antigens are not always available, and relying solely on one strategy severely limits the pool of optimal antigen pairs. To overcome these limitations, we developed an inside-out strategy that combines logic-gated CARs with PC-CARs, enabling recognition of both intracellular and extracellular proteins.

synNotch System

Synthetic Notch (synNotch) is an engineered receptor system that enables regulation of gene expression upon antigen recognition. The first CAR binds to a primary antigen and activates the Notch receptor, which then cleaves and releases a transcription factor. With the release of the transcription factor, expression of a second CAR is induced. This serial activation serves as an AND gate (If-Then) and activates T-cells only if both target antigens are present.

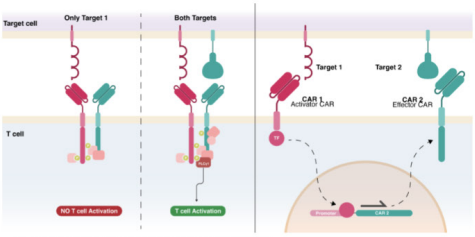


Figure A: Visual explanation of the SynNotch System

Results

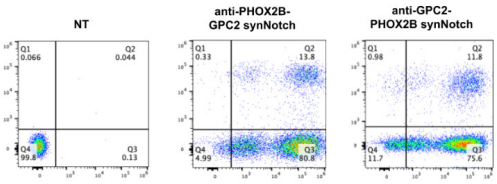


Figure B: Flow cytometry analysis showing NT, anti-PHOX2B-GPC2 synNotch, and anti-GPC2-PHOX2B synNotch expressions. The Non-transduced (NT) cells did not have any CAR expression while cells that were co-transduced with both constructs had both expressions at the same time proving that the inside-out strategy works.

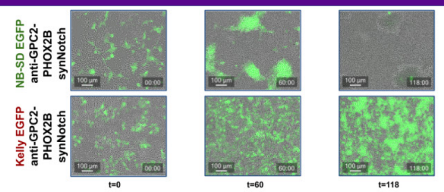


Figure C: Exemplary Images from the killing assay showing target (NB-SD EGFP, HLA-A23:01) and non-target (Kelly EGFP, HLA-A02:01) cells at 0, 60, and 118 hours. EGFP signal form NB-SD target cells disappear over-time, indicating target cell killing,

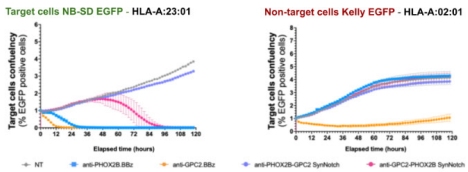


Figure D: Inside-out (IO)-SynNotch CAR T cells killing assay on target and non-target cells. IO-anti-GPC2-PHOX2B SynNotch CAR T cells kills specifically NB-SD target cells while sparing non-target Kelly cells, which do not present the correct HLA.

Future Work

For the future work, different antigen pairs can be identified. Additionally, the in-vitro findings can be evaluated in vivo. At the end, it can be followed by preclinical optimization for potential clinical translation

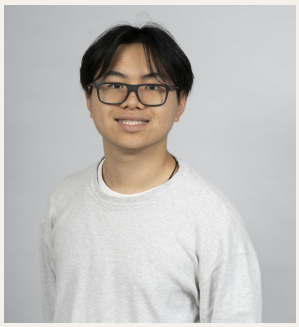
Works Cited

[1] Yarmarkovich, M., Marshall, Q.F., Warrington, J.M. et al. Targeting of intracellular oncoproteins with peptide-centric CARs. Nature 623, 820–827 (2023). https://doi.org/10.1038/s41586-023-06706-0
[2] Nolan-Stevaux O and Smith R (2024) Logic-gated and contextual control of immunotherapy for solid tumors: contrasting multi-specific T cell engagers and CAR-T cell therapies. Front. Immunol. 15:1490911. doi: 10.3389/fimmu.2024.1490911

Acknowledgement

I would like to thank NYU Tandon School of Engineering's Office of Undergraduate Academics for the funding of the project. Also, special thanks to Dr.Michele Palamenghi for his guidance and Dr.Mark Yarmarkovich for the opportunity.

Engineering Health



Cayden Auyeung
Biomolecular Science, B.S.
Spring 2028
NYU Tandon



David Kedem
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

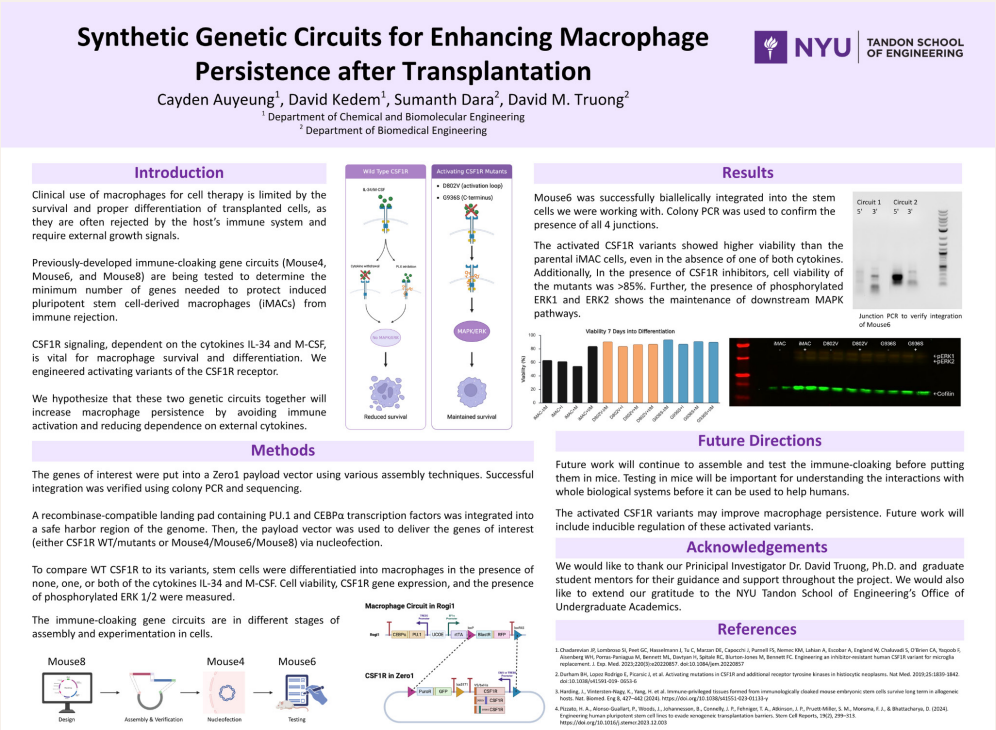
Engineering Universal Donor Macrophages from Stem Cells

Macrophages have great potential as cell therapies because they can enter damaged tissues, remove harmful material, and help organize tissue repair. However, clinical use is limited because transplanted cells are often rejected by the immune system and require external growth signals to survive. Our goal is to engineer macrophages to overcome both of these challenges and serve as universal donor cells.

To do this, we are engineering stem cells with synthetic genetic circuits. We are testing previously-developed immune-cloaking systems we refer to as Mouse4, Mouse6, and Mouse8, each named based on the number of overexpressed immunomodulatory transgenes that protect transplanted cells from immune rejection. We want to determine the minimum number of genes needed to accomplish this task. We are also incorporating an inducible PU.1 and CEBPα transcription factor circuit to differentiate these cells into macrophages. Finally, we are also introducing engineered variants of CSF1R, a receptor that regulates macrophage development, to promote survival even when external cytokines are not present.

Preliminary results show that the engineered CSF1R variants maintain viability during macrophage differentiation without the cytokines normally required for survival. Ongoing experimentation is evaluating whether the simplified immune-cloaking circuits provide protection like in Mouse8 and whether immune-cloaked stem cells can be successfully differentiated into macrophages.

Together, these experiments lay the foundation for increasing the persistence of engineered macrophages through evading immune rejection, ensuring proper differentiation, and decreasing their dependence on external signaling molecules.



Engineering Health



Celia Meunier
Biomolecular Science, B.S.
Spring 2027
NYU Tandon
Thompson-Bartlett Fellow



Natalie Nylaende
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon

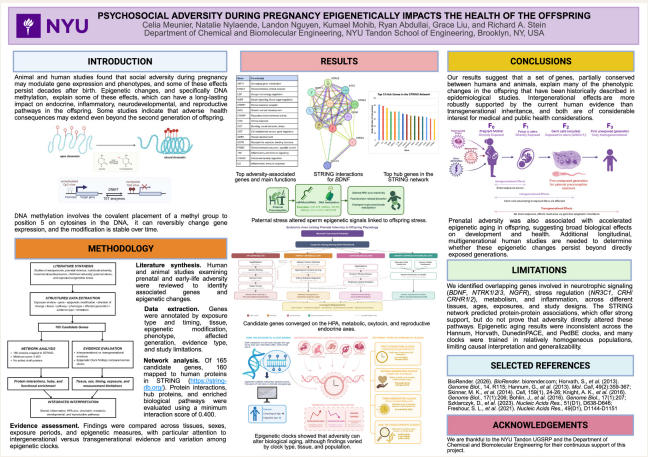


Landon Nguyen
Biomolecular Science, B.S.
Fall 2029
NYU Tandon

Ryan Abdullai
Major
Summer 2026
NYU Tandon

Gestational Adversity Epigenetically Impacts Multiple Signaling Pathways in the Offspring

Data from animal and human studies have increasingly found that social adversity during pregnancy may modulate gene expression and phenotypes, and some of these effects are transmitted to the offspring. Epigenetic changes, and specifically DNA methylation, were identified as a mechanism that explains these effects, which can have a long-lasting impact on endocrine, inflammatory, neurodevelopmental, and reproductive pathways in the offspring. Methods: We interrogated publicly available databases of genes that changed their DNA methylation status and expression levels in the offspring of women who experienced gestational adversity, including survivors of genocide, famine, prenatal violence, and emotional stress. Exposure time, tissue type, the impacted generation, the type of evidence, and study limitations were noted for every candidate gene. Gene candidates with the highest confidence levels were analyzed using the STRING protein-interaction database to predict functional association networks and the affected cellular pathways, and the Drug-Gene Interaction Database (DGIdb) was used to identify potential therapeutic leads that could reverse epigenetic signatures of adversity. In parallel, we performed an analysis of genes that are epigenetically modulated in the offspring of pregnant rodents undergoing adversity. Results: We identified 876 significant protein-protein interactions, of which 278 were of high confidence and 125 of very high confidence. GAPDH, IL6, TNF, BDNF, HDAC1, IL1B, POU5F1, NR3C1, HSP90AA1, and SIRT1 were among the most prominent hub proteins. Reactions to endogenous stimuli, hormones, lipids, oxidative chemicals, and reproductive, developmental, and metabolic activities were all over-represented by functional enrichment, and results differed according to tissue, sex, exposure type, developmental time, and epigenetic measure. These results point to a convergence of chromatin alteration, inflammation, reproductive programming, and the control of the hypothalamic-pituitary-adrenal axis that are modulated in the offspring as a result of gestational adversity. Conclusion: Our results indicate that a set of genes, partially conserved between humans and animals, explain phenotypic changes in the offspring, explain many of the phenotypes that have been historically described in epidemiological studies. Intergenerational effects are more robustly supported by the current human evidence than transgenerational inheritance, and epigenetic age acceleration appears to be one of the phenotypic consequences that integrates the various effects on organs and systems. To understand how extreme stress gets biologically ingrained and whether these alterations endure beyond directly exposed generations, further long-term, multigenerational human studies are required.





Cristiana Pintea
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

Specificity assessment of pathogens targeting aptamers using Biolayer Interferometry

According to the CDC, one in five Americans has a sexually transmitted disease (STD), and the World Health Organization estimates that there are about 374 million new cases each year around the globe, making early detection critical for treating patients and preventing outbreaks. Biosensors have proven extremely valuable in this effort, as they are non-invasive, provide rapid results and do not require a clinical setting to operate. Some of these biosensors work by converting the binding between analytes and specific recognition molecules, or bioreceptors, into readable electrical signals. Bioreceptors are designed to bind to unique targets, such as proteins and small molecules, and are used to identify the presence of disease biomarkers in bodily fluids. Antibodies are the traditional biorecognition molecules used in biosensing due to their high specificity for target antigens, and are the basis for widely used immunological assays such as ELISA. However, antibody production is costly, limiting their application in low-cost biosensing. Aptamers, single-stranded DNA or RNA sequences offer a promising alternative, as they are chemically synthesized to have high affinity and specificity for a wide range of targets. In this work, we used two aptamers selected for AP65, a membrane protein from Trichomonas vaginalis, and OmpA, a membrane protein for Chlamydia trachomatis, both common and frequently underdiagnosed STIs. Through Biolayer Interferometry, BLI, we measured the binding affinity between the selected aptamers and the target proteins. Each aptamer was tested both against its target protein and against the non-target protein to assess specificity and cross-reactivity. This work lays the groundwork for using aptamers as recognition receptors for the detection of T. vaginalis and C. trachomatis. Future experiments will focus on integrating these probes onto a biosensing platform and extending this characterization approach to additional STI biomarkers. These efforts will lead towards a multiplexed diagnostic platform for STIs.

NYU TANDON SCHOOL OF ENGINEERING

Specificity assessment of pathogen - targeting aptamers using Biolayer Interferometry for biosensor development

Cristiana Pintea[1], Tausif Tamim[1], Elisa Riedo[1], Paola Bardetti[1]
[1]NYU Tandon School of Engineering, Department of Chemical and Biomolecular Engineering

ABSTRACT

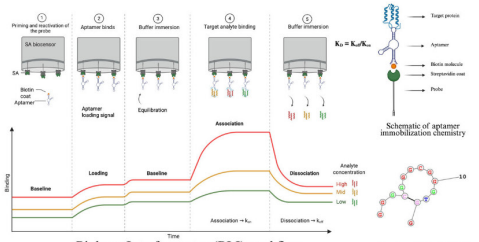
Trichomoniasis, caused by the parasite *Trichomonas vaginalis*, is one of the most prevalent STDs in the United States, affecting roughly 3% of the population. Chlamydia, caused by *Chlamydia trachomatis*, is similarly widespread and frequently overlooked. With up to 70% of Trichomonas infections presenting asymptomatically, many cases go undiagnosed, underscoring the need for accessible and rapid detection methods. Biosensors, such as the glucose monitor patch, offer a promising solution, as they provide non-invasive, real-time detection of disease biomarkers and they do not require clinical setting. Biosensors require a recognition molecule, a target analyte, a transducer, and a signal processor. Aptamers are single-stranded DNA or RNA sequences engineered to bind to various targets, such as proteins or small molecules, with high affinity and specificity. Due to these characteristics, aptamers are promising tools for disease detection. In this work, we characterized the affinity and specificity of two aptamers selected against two target proteins: AP65, a surface adhesion protein expressed by *Trichomonas vaginalis*, and OmpA, an outer membrane protein expressed by *Chlamydia trachomatis*. We assessed the affinity (K_D) and kinetics between the aptamers and the target proteins using Biolayer Interferometry (BLI). Moreover, each aptamer was also tested against the non-target protein to evaluate cross-reactivity, which is critical to avoid non-specific binding during sensing. These experiments lay the groundwork for their use in a rapid, point-of-care, multiplexed electrical and electrochemical biosensing platform.

OBJECTIVES

- Characterize binding affinity and kinetics of the AP65 A1 aptamer against *Trichomonas vaginalis* AP65 protein via BLI
- Characterize binding affinity and kinetics of the R10-16 aptamer against *Chlamydia trachomatis* OmpA protein via BLI
- Assess cross-reactivity between aptamer-protein pairs to evaluate specificity for a multiplexed biosensing platform

MATERIAL & METHODS

Aptamers selected against AP65 and OmpA proteins were synthesized with a 5' biotin tag and prepared at a 200nM concentration for immobilization on streptavidin-coated biosensor tips. Protein concentration ranges were selected based on literature-reported K_D values, ensuring the tested range would capture both saturating and sub-saturating binding conditions.

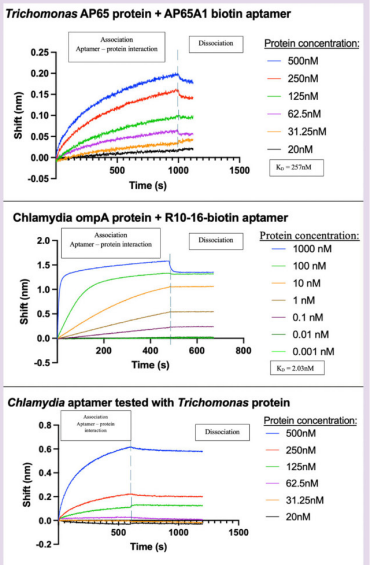


Biolayer Interferometry (BLI) workflow

Schematic of aptamer immobilization chemistry

Structure of the AP65 A1 aptamer

RESULTS



BLI sensorgrams showing concentration-dependent binding of AP65 A1 to *Trichomonas* AP65 protein (top), R10-16 to *Chlamydia* OmpA protein (middle), and R10-16 cross-reactivity against AP65 protein (bottom)

CONCLUSIONS

We successfully characterized the binding and kinetics of the AP65 aptamer-AP65 protein and *Chlamydia* aptamer-OmpA protein interactions. Preliminary experiments showed signs of cross-talk when testing the *Trichomonas* AP65 protein on the R10-16 aptamer. Future work will focus on evaluating the cross-reactivity, as assay conditions may influence specificity. Further investigation into the structural and sequence-level similarity between AP65 and OmpA may help clarify the source of the cross-talk, despite their differing functions as an adhesion protein and a porin, respectively. Additional cross-reactivity testing against aptamers for other STI targets will also be pursued to build a more complete specificity profile before integrating into the multiplexed platform.

FUTURE DIRECTIONS

- Characterizing binding using Surface Plasmon Resonance (SPR), which is more specific and better mimics the surface of the biosensor
- Test longer aptamer sequence variants to assess whether extended length improves binding affinity or specificity
- Refine tested protein concentration ranges based on observed binding curves to better resolve affinity and specificity
- Optimize aptamer sequence, as its G-rich composition may promote G-quadruplex formation, potentially contributing to the observed cross-reactivity
- Refine assay buffer conditions to optimize sensor signal quality and reduce background noise

ACKNOWLEDGEMENT

We thank the NYU Undergraduate Summer Research Program for supporting this work, Prof. Paola Bardetti and Prof. Elisa Riedo for their mentorship and guidance throughout this project.

PicoForce Lab
www.picoforcelab.org

REFERENCES

[1] Akbari Nakhjavani, S., Khoshbin, H., Ghoreis, S., & Omidfar, K. (2024). Advances in Biosensor technologies for infectious diseases detection. TrAC Trends in Analytical Chemistry, 180, 117970. https://doi.org/10.1016/j.trac.2024.117970

[2] CDC - (2024). Sexually transmitted infections prevalence, incidence, and cost estimates in the united states. Sexually Transmitted Infections (STIs). https://www.cdc.gov/sti-data-research/prevalence-estimates/index.html

[3] Mehta, S. (2011). Trichomoniasis. Infectious Disease Clinics of North America, 25(4), 739–749. https://doi.org/10.1016/j.idc.2011.09.003

[4] Tretiak, M., Ai, H., Kolson, Warren N., Tilburu, R., & Alexander, N. (2024). State-of-care biosensors for infectious disease diagnosis. Recent updates and future directions. Biosensors and Bioelectronics: X, 19, 100494–100494. https://doi.org/10.1016/j.biosx.2024.100494

[5] Nadia, N., Tiwari, A., Pandya, A., & Tripathi, D. (2025). Nanoparticle-biosensors for infectious disease surveillance: Innovations, challenges, and global health impact (Vol. 313, pp. 125–478). Academic Press. https://doi.org/10.1016/bs.pbr.2025.02.010

[6] Zhou, Z., Yu, Y., Wang, M., Li, J., Zhang, Z., Liu, J., Wu, X., Lu, A., Wang, G., & Zhang, G. (2017). Recent Advances in Saliva Technology and Salivaomic Applications in Biomedicine. International Journal of Analytical Bioscience, 5(1), 21-27. https://doi.org/10.12691/ijab-5-1-4

Created with BioRender Poster Builder

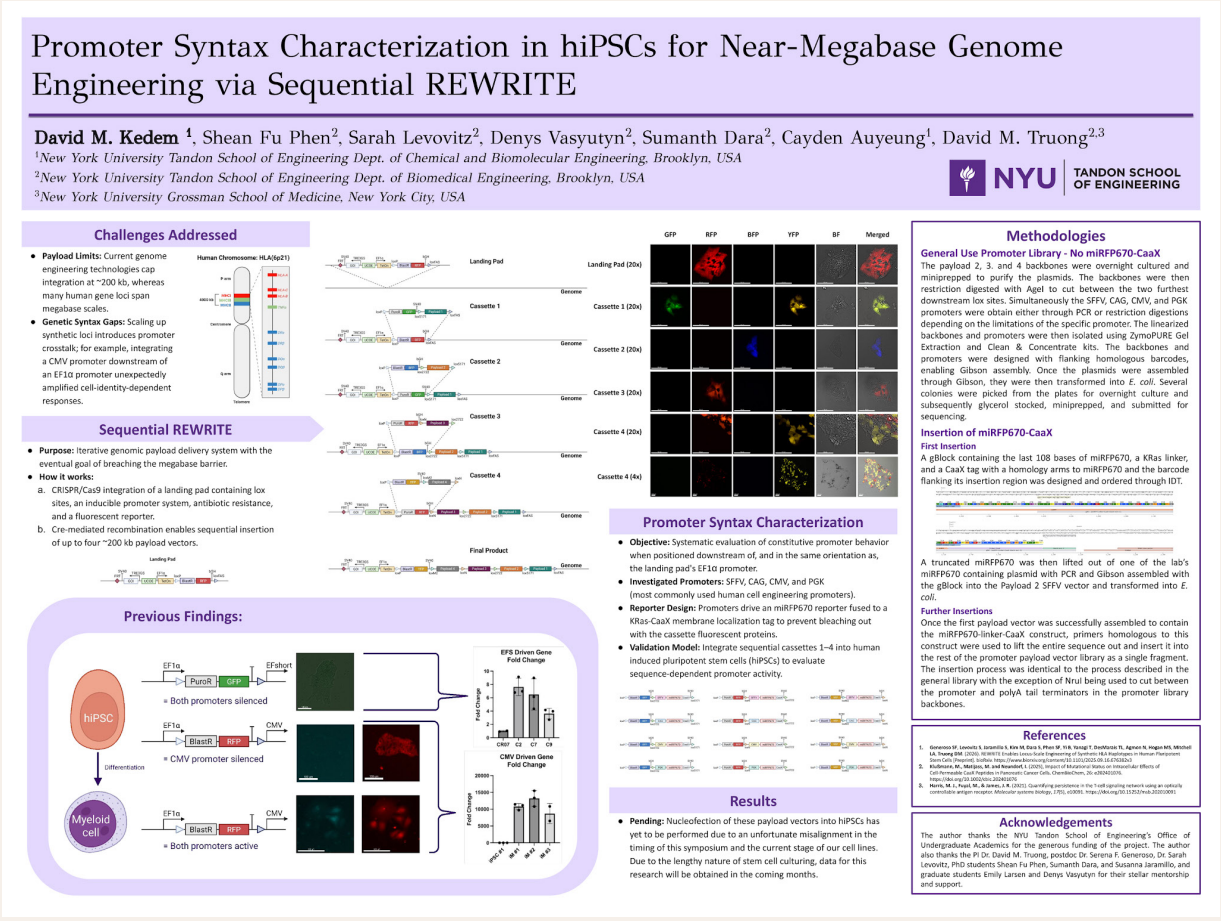
Engineering Health

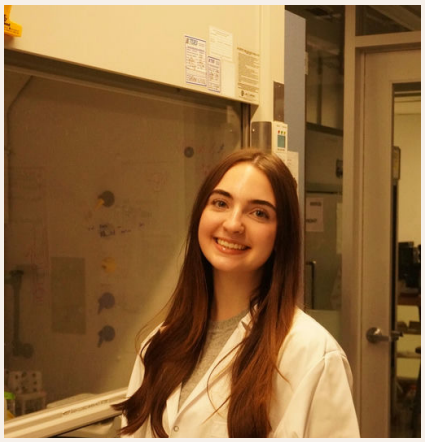


David Kedem
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

Promoter Syntax Characterization in hiPSCs for Near-Megabase Genome Engineering via Sequential REWRITE

Current human genome engineering technologies are limited to integrating payloads of up to ~200 kb, whereas many human gene loci span megabases. To address this, the Truong Lab is developing Sequential Recombinase Writing by Repeatable Integrations and Tag Excision (Sequential REWRITE), a platform designed for near-megabase genomic modifications. Sequential REWRITE begins with CRISPR/Cas9-mediated integration of a genomic landing pad containing lox sites, an inducible promoter system, antibiotic resistance, and a fluorescent reporter. Cre-mediated recombination then enables sequential integration of currently up to four ~200 kb payload vectors, reaching a total edit of ~800 kb. Each cycle exchanges selection markers to allow the identification and isolation of correctly engineered cells, as well as enable resistance gene recycling. However, genetic syntax at this scale remains poorly understood; specifically, how multiple constitutively active promoters interact in series. Previously, the Truong Lab integrated the CMV promoter downstream of the EF1α promoter and found that its cell identity dependent response was greatly enhanced. To enable the predictable engineering of genetic circuits, this challenge must be addressed. Therefore, we are investigating how SFFV, CAG, CMV and PGK—the most frequently used promoters in human cell engineering—function within payload vectors when integrated downstream of, and in the same orientation as, the landing pad's constitutively active EF1α. To evaluate promoter activity, we designed an miRFP670 reporter fused to a KRas-CaaX membrane localization tag and placed it under the control of the investigated promoters. Sequential cassettes 1, 2, 3, and 4 were assembled and integrated into human induced pluripotent stem cells (hiPSCs) to systematically evaluate promoter performance. These will eventually provide insights establishing foundational design rules for megabase-scale synthetic loci, driving further optimization of the Sequential REWRITE platform in clinically relevant hiPSCs.

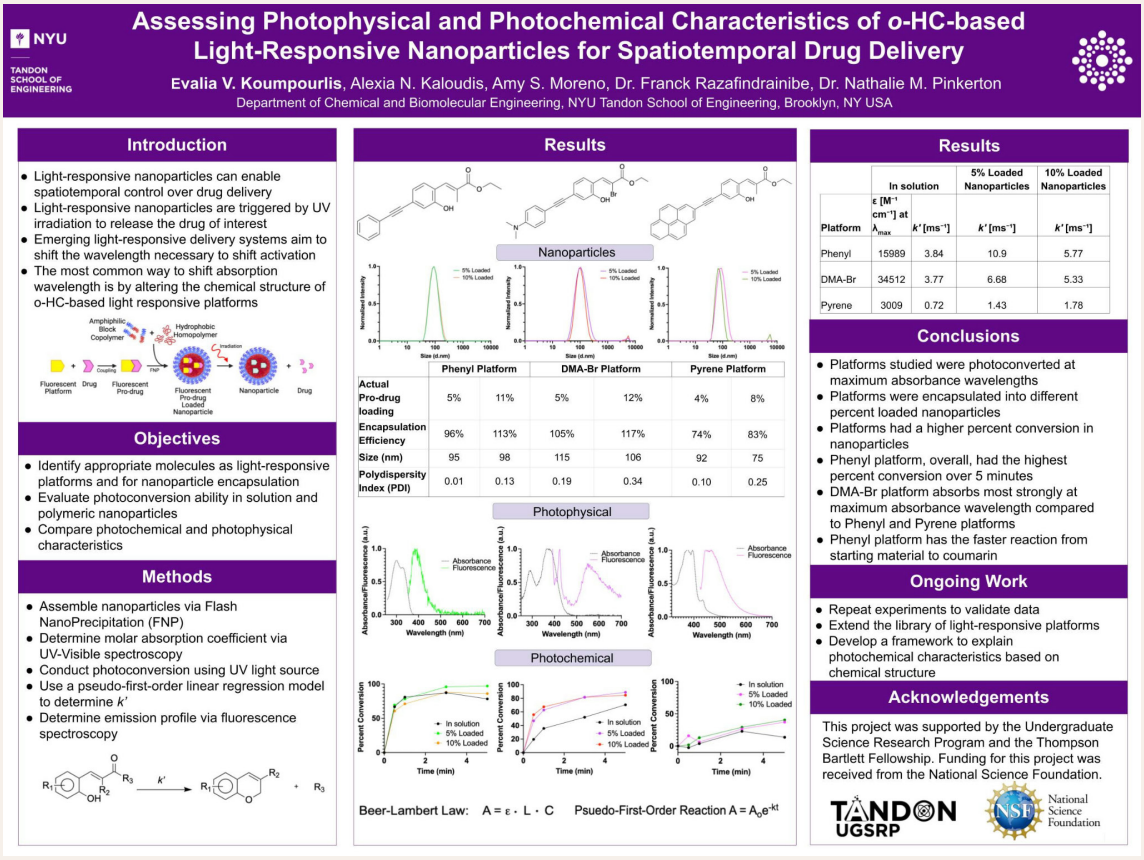




Evalia Koumpourlis
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon
Thompson-Bartlett Fellow

Nanoparticles for Drug Delivery with Enhanced Spatiotemporal Control via Light

Chronic pain and cancer affect around 20% (1) of the population, emerging drug delivery systems seek to improve pain management through nanoparticles. Nanoparticles are vehicles to deliver drugs to various parts of the body. This research presents a new light responsive drug delivery system with enhanced spatiotemporal control for improved delivery. Four ortho-Hydroxycinnamoyl (o-HC) based light responsive molecules were synthesized and analyzed for the purpose of this research. The molecules analyzed differed in functionality with regards to the groups para to the hydroxyl group and on the alpha position on the unsaturated bond. By changing the functional group para to the hydroxyl group and on the alpha position, we saw that maximum absorbance wavelength increased displaying a red shift. Furthermore, this change saw an increase in maximum emission wavelength. The molecules analyzed displayed a stokes shift. The molecules were UV irradiated, in solution, near their maximum absorption wavelengths to observe the percent of conversion. Additionally, the molecules were encapsulated at different loading percentages via Flash Nano-Precipitation and UV irradiated to note percent conversion. Across the molecules, the loading capacity threshold was around 10%. The molecule with CN on the alpha carbon of the unsaturated bond does not photoconvert. Encapsulation in nanoparticles contributes to higher percent conversions across the remaining molecules in comparison to in solution which aligns with previous data published from the Pinkerton lab (2). The molecule with a phenyl group para to the hydroxyl group converted the fastest overall. The reaction rate of photoisomerization, k prime, was calculated to determine quantitatively which molecule is the fastest. The molecule with the phenyl group in solution, had the fastest reaction rate. This research aids in making a more hydrophobic drug that enables delivery at specific times and locations as intended for. This project is impactful because these molecules provide more spatiotemporal control of the drug delivery process. Moreover, these molecules are instrumental in shifting wavesources to more biocompatible wavelengths.



Engineering Health



Fanni Miettinen
Mechanical Engineering, B.S.
Spring 2029
NYU Tandon
Thompson-Bartlett Fellow



Chiara Pinci
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon



David Chauca
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon



Aiden Li
Electrical Engineering and Computer Engineering , B.S.
Fall 2027
NYU Tandon



Mark Wang
Electrical Engineering and Computer Engineering (dual major), B.S.
Spring 2027
NYU Tandon



Vicente Correia
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon



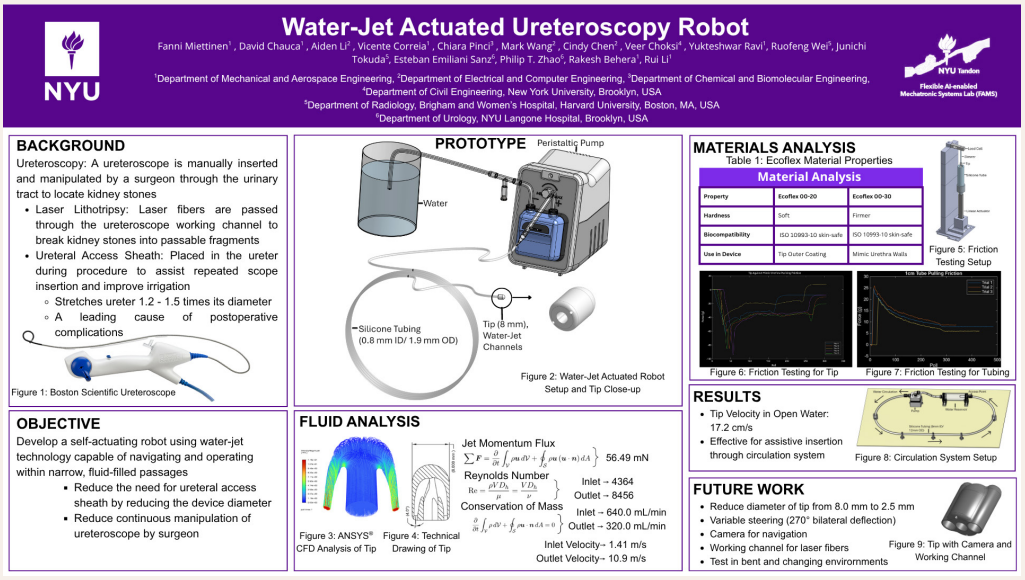
Veer Rushin Choksi
Civil Engineering, B.S.
Spring 2027
NYU Tandon



Cindy Chen
Electrical Engineering and Computer Engineering (dual major), B.S.
Spring 2029
NYU Tandon

Water-Jet Actuated Ureteroscopy Robot

Conventional flexible ureteroscopy relies on continuous manual manipulation by the surgeon. This often benefits from the use of a ureteral access sheath for repeated scope insertion and improved irrigation. However, access sheath placement may increase the risk of ureteral injury. To address these limitations, we developed a soft, tethered water-jet actuating capsular robot able to navigate through the narrow, fluid-filled passages of the urinary tract with minimal external manipulation. A peristaltic pump supplies water through soft flexible silicone tube to a 5 mm tip with two adjacent water channels which generate forward propulsion. The robot can achieve fast locomotion, measuring 10.75 cm/s in open water environments. Fluid dynamics modeling of the robot has been tested in ANSYS®. Experimental testing has two components: a test rig to measure the amount of resistive force the robot needs to overcome and a water-based circulation system used to mimic flow of irrigation and urine during procedures. Both ensure the robot's capability to maneuver within the enclosed passages of the urinary tract. With reduced size and increased self-locomotion capabilities, the robot has the potential to minimize patient discomfort and manual contortion performed by physicians during urinary tract surgeries.



Engineering Health



Fei Wu
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

Lymph Node-on-a-chip

Lymph Node-on-a-chip technology offers an advanced platform for recreating physiological tissue microenvironments in vitro. In this project, a three-channel lymph node-on-a-chip microfluidic device was used to model key features of the human lymph node microenvironment. A hydrogel-based extracellular matrix was prepared and loaded into the device to support three-dimensional cell culture and cellular organization within the microfluidic channels.

Fluorescence imaging using multiple wavelengths was performed to characterize immune cell marker expression and spatial organization within the microfluidic device. Microscopy images were analyzed through manual cell identification and counting, as well as assessment of fluorescence signals across different imaging channels. These analyses were used to evaluate cell distribution, marker expression, and cellular localization within the lymph node-on-a-chip platform.

This work contributes to the development and characterization of an in vitro lymph node model and demonstrates the application of hydrogel preparation, microfluidic cell culture, multi-channel fluorescence imaging, and quantitative image analysis in organ-on-a-chip research.

NYU TANDON SCHOOL OF ENGINEERING

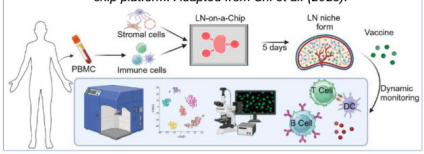
Evaluation of Vaccine-Associated CD25 Expression in CD4⁺ T Cells Using a Human Lymph Node-on-a-Chip Model

Fei Wu
Department of Chemical and Biomolecular Engineering
New York University Tandon School of Engineering

Background

Human lymph node-on-a-chip (LN-on-a-chip) models provide a biomimetic platform for studying adaptive immune responses in vitro [1]. Compared with conventional cell culture, this platform better recreates the human lymph node microenvironment, allowing dynamic evaluation of immune activation following vaccination. Within this platform, T-cell activation can be evaluated using fluorescence microscopy. CD25 is an early activation marker expressed on activated T cells and serves as a quantitative indicator of immune activation [2]. Therefore, measuring CD25 fluorescence intensity provides a useful approach for assessing vaccine-associated immune responses.

Figure 1: Schematic illustration of the human lymph node-on-a-chip platform. Adapted from Shi et al. (2025).

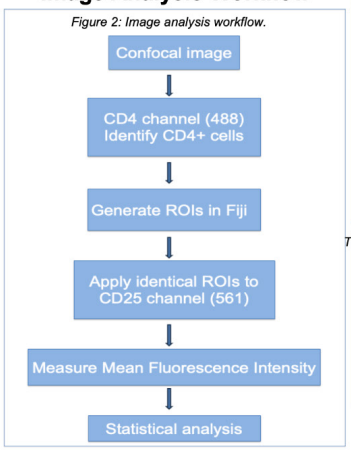


Objective

To evaluate vaccine-associated changes in CD25 expression in CD4⁺ T cells using confocal microscopy and quantitative image analysis.

Image Analysis Workflow

Figure 2: Image analysis workflow.



Methods

Human lymph node-on-a-chip samples were cultured under control or vaccine-treated conditions. Then confocal images were collected on Days 5 and 9. CD4⁺ T cells were identified, and the regions of interest were generated in Fiji and transferred to the corresponding CD25 images. CD25 mean fluorescence intensity was quantified and analyzed using two-way ANOVA.

Results

Figure 3: Representative confocal images.

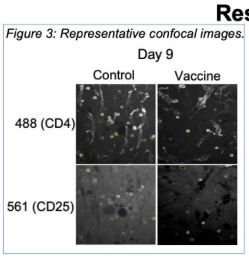


Figure 4. CD25 fluorescence intensity in CD4⁺ T cells (mean ± SEM, n = 3).

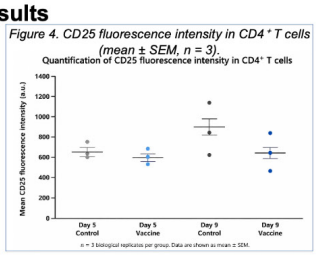


Table 1. Two-way ANOVA results for CD25 fluorescence intensity.

Two-way ANOVA	F	p
Day	2.097	0.186
Treatment	2.611	0.145
Day × Treatment	0.965	0.355

- CD25 fluorescence intensity tended to be higher in the Day 9 control group than in the other groups.
- Two-way ANOVA revealed no significant effects of day, treatment, or their interaction ($p > 0.05$).
- These findings suggest that vaccine treatment did not significantly alter CD25 expression under the current experimental conditions.

Conclusion

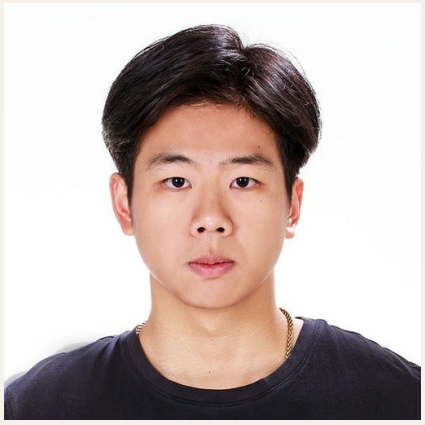
CD25 fluorescence intensity was successfully quantified in CD4⁺ T cells using a Fiji-based ROI analysis workflow. Although no statistically significant differences were detected between groups, this workflow provides a reproducible approach for evaluating immune activation in the human LN-on-a-chip model.
Future studies with larger sample sizes and additional activation markers will further evaluate vaccine-induced immune responses in the LN-on-a-chip platform.

Works Cited

[1] Shi F, Liu D, Song Y, et al. A Human Lymph Node-on-a-Chip for Personalized Evaluation of Vaccine Immunogenicity. Research Square. 2026. doi:10.21203/rs.3.rs-8951191/v1.
[2] Ptacin JL, Ma L, Caffaro CE, et al. A CD25-biased interleukin-2 for autoimmune therapy engineered via a semi-synthetic organism. Communications Medicine. 2024;4:58.
[3] Schindelin J, Arganda-Carreras I, Frise E, et al. Fiji: an open-source platform for biological-image analysis. Nature Methods. 2012;9:676–682.

Acknowledgement

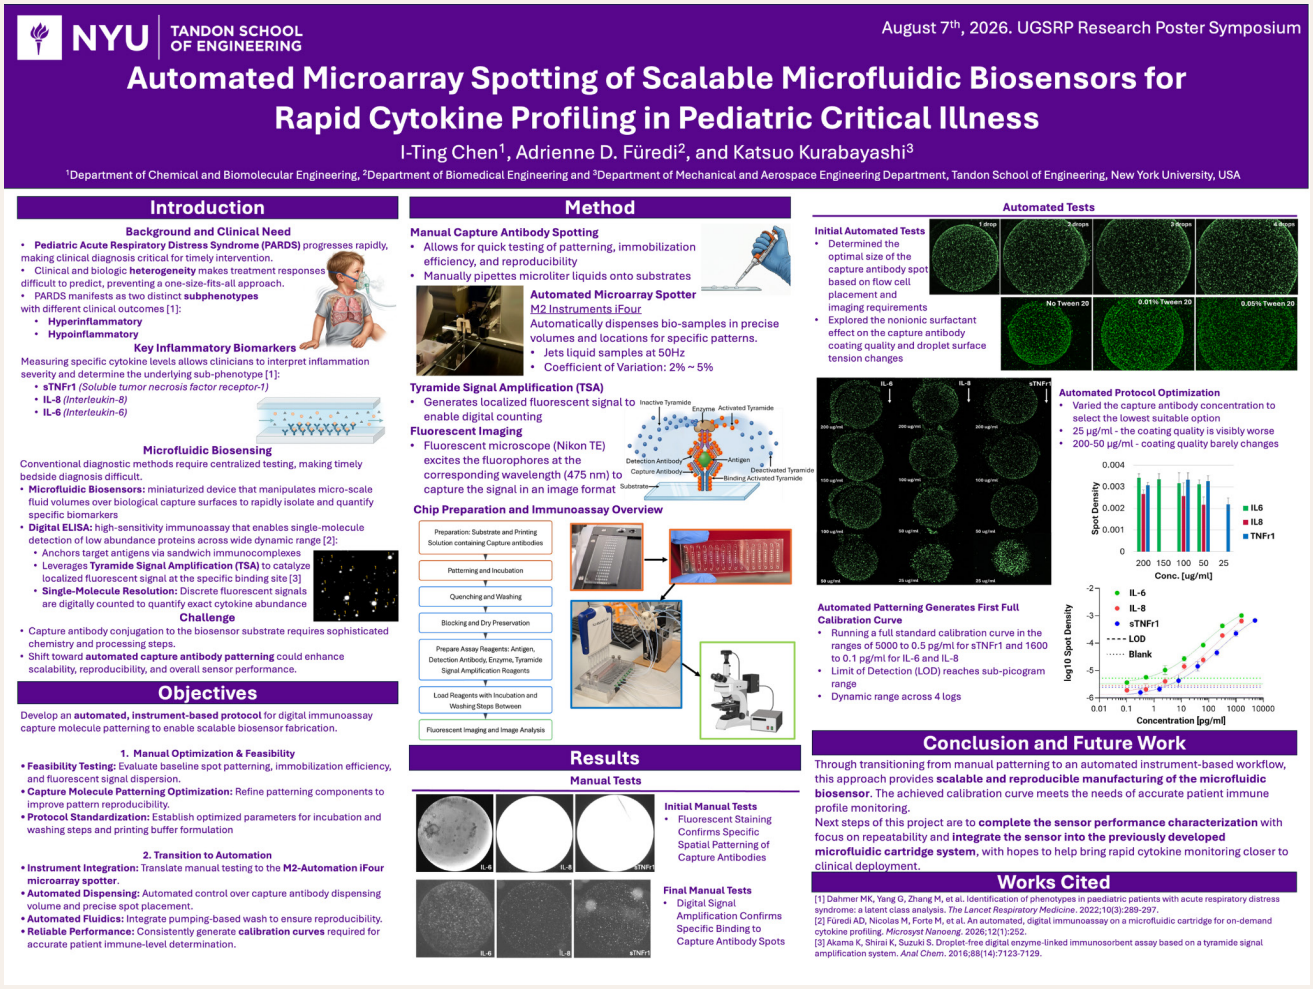
The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.



Cody Chen
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon

Automated Microarray Spotting of Scalable Microfluidic Biosensors for Rapid Cytokine Profiling in Pediatric Critical Illness

Pediatric Intensive Care Unit (ICU) patients require urgent diagnosis of immune system profiles to guide clinical intervention. This project contributed to the advancement of an automated microfluidic digital immunoassay system, engineered to promptly provide quantitative data based on the levels of key inflammatory biomarkers - Interleukin-6 (IL-6), Interleukin-8 (IL-8), and Tumor Necrosis Factor Receptor 1 (TNFr1). This development phase focused on establishing novel manufacturing methods of the system's biosensor chips by utilizing automated microarray spotter technology. The development progressed through two stages. Manual arraying tests were first conducted to define baseline patterning feasibility and immobilization efficiency of the capture molecules. These tests allowed for inspection of the immune capture dynamics and patterning reproducibility. Through the manual tests, we determined a washing protocol and printing buffer formulation optimal for the biosensor. The manual tests were then translated to a microarray spotter instrument to automate the deposition routines; this produced precise and consistent spot morphology. Via characterization experiments, we optimized spot size to fit up to an 8-plex panel per channel without droplets merging and reduced capture antibody concentration from 200 µg/mL to 50 µg/mL while maintaining uniform coating. This established a standardized protocol to consistently generate calibration curves for immune-level determination. Through transitioning from manual patterning to an automated instrument-based workflow, this approach provides scalable and reproducible manufacturing of the microfluidic biosensor, which may help bring rapid cytokine monitoring closer to clinical deployment.



Engineering Health

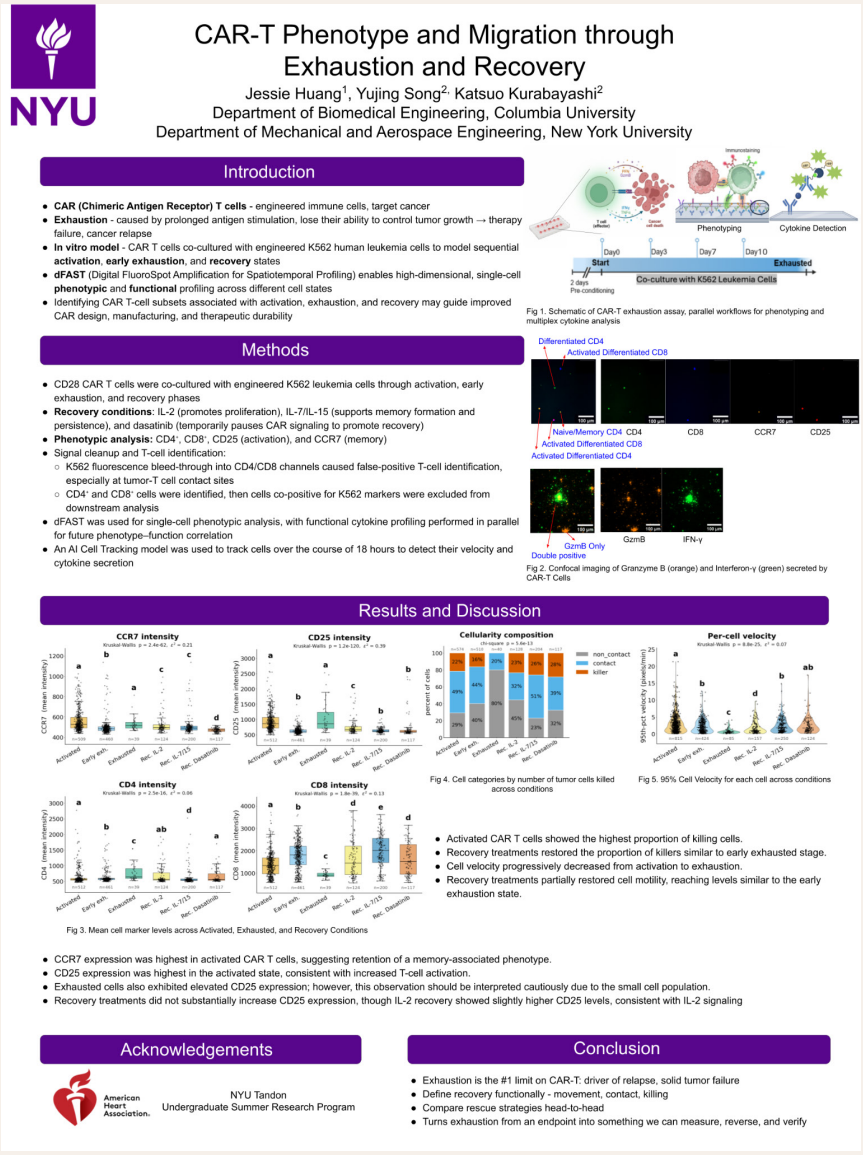


Jessie Huang
Biomedical Engineering,B.S.
Spring 2027
Columbia University
AHA Fellow

Single-Cell Characterization of CAR-T Cell Exhaustion and Recovery Through Phenotypic Profiling and AI-Assisted Tracking

Chimeric antigen receptor T (CAR-T) cell therapy has demonstrated remarkable success in treating hematological malignancies. However, T-cell exhaustion remains a major challenge that limits therapeutic efficacy and persistence. Understanding how exhaustion influences CAR-T cell behavior at the single-cell level is essential for improving treatment outcomes. This study investigates the interactions between CAR-T cells and K562 tumor cells by integrating live-cell imaging, computational cell tracking, phenotypic profiling, and cytokine secretion analysis to characterize functional changes associated with activated and exhausted CAR-T states.

Live-cell imaging datasets of CAR-T and K562 co-cultures are analyzed using an AI-assisted cell detection and tracking pipeline to reconstruct individual cell trajectories over time. The tracking outputs are integrated with phenotypic measurements, including CD25 and CCR7 expression, and functional classifications such as serial killer and single killer CAR-T cells. Cytokine secretion and migration metrics are further incorporated to examine relationships between immune phenotype, cell motility, and cytotoxic activity. Ground-truth annotations are used to evaluate tracking performance and support downstream biological analyses.



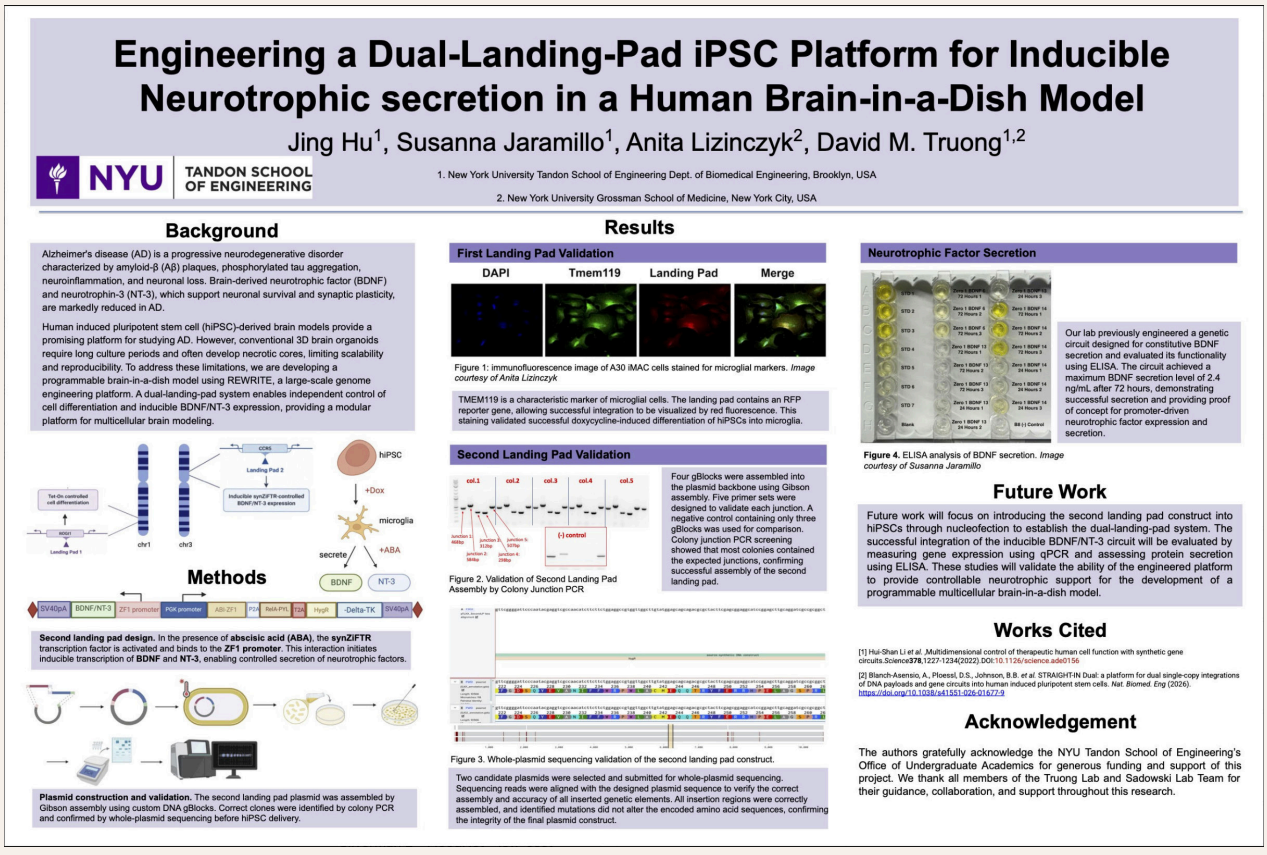
The study aims to provide quantitative comparisons of phenotype, migration dynamics, killing behavior, and cytokine secretion across activated, early exhausted, and recovery-stage CAR-T cell states. By combining computational image analysis with single-cell immunophenotyping, this work seeks to establish a comprehensive framework for investigating CAR-T exhaustion and identifying phenotypic and functional biomarkers that may inform future CAR-T optimization and therapeutic development.

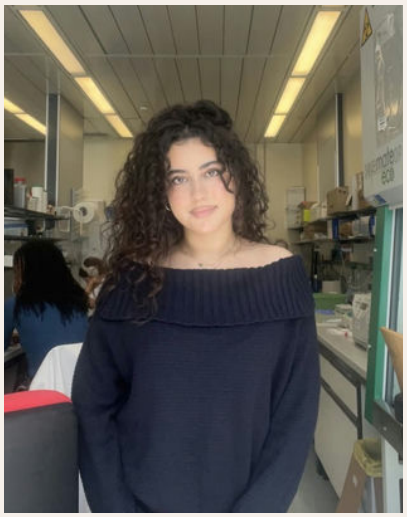


Jing Hu
Biomolecular Science, B.S.
Spring 2027
NYU Tandon

Engineering a Dual-Landing-Pad iPSC Platform for Inducible Neurotrophic Secretion in a Human Brain-in-a-Dish Model

Alzheimer's disease (AD) is a neurodegenerative disorder characterized by amyloid-β (Aβ) plaque accumulation, intracellular phosphorylated tau aggregation, lipid dysregulation, innate and adaptive immune cell dysregulation, and progressive neuronal loss. Neurotrophic factors, including brain derived neurotrophic factor (BDNF) and neurotrophin-3 (NT-3), play essential roles in neuronal survival and synaptic plasticity, yet their expression is markedly reduced in patients with AD. Although 3D brain organoids are considered the gold standard for human brain modeling, they have several limitations, including the formation of necrotic cores as well as prolonged culture times. To address these limitations, we aim to develop an inducible genetic circuit that enables controllable expression of BDNF and NT-3 in human induced pluripotent stem cells (hiPSCs). This approach eliminates the need for supplementation of culture media with neurotrophic factors while providing sustained trophic support to promote neuronal survival, maturation, and network development. Using REWRITE, a platform for large-scale genome engineering developed by our lab, we engineered hiPSCs with circuits for forward programming into multiple brain cell types, enabling the establishment of a multicellular brain-in-a-dish model composed of four brain cell types including microglia capable of constitutive BDNF and NT-3 secretion. To enable inducible secretion of neurotrophic factors while preserving the differentiation capacity of engineered iPSCs, we designed a dual-landing-pad system that integrates programmable genetic circuits into two independent genomic loci. The existing landing pad at the ROGI1 locus directs Tet-On-mediated differentiation of hiPSCs into induced microglia (iMs), while a second landing pad is designed for targeted integration at the CCR5 safe harbor locus to enable synZiFTR-regulated inducible expression of BDNF and NT-3. Using Gibson assembly, Golden Gate assembly, and yeast assembly, we constructed a series of plasmids encoding the inducible neurotrophic factor circuit. Future work will introduce the second landing pad into the CCR5 locus, and functional expression and secretion of BDNF and NT-3 will be evaluated by qPCR and ELISA. Together, this dual-landing-pad platform provides a modular strategy for independently regulating cell fate and neurotrophic factor expression. This system establishes a foundation for a programmable, self-sustaining human brain-in-a-dish model to investigate AD mechanisms.



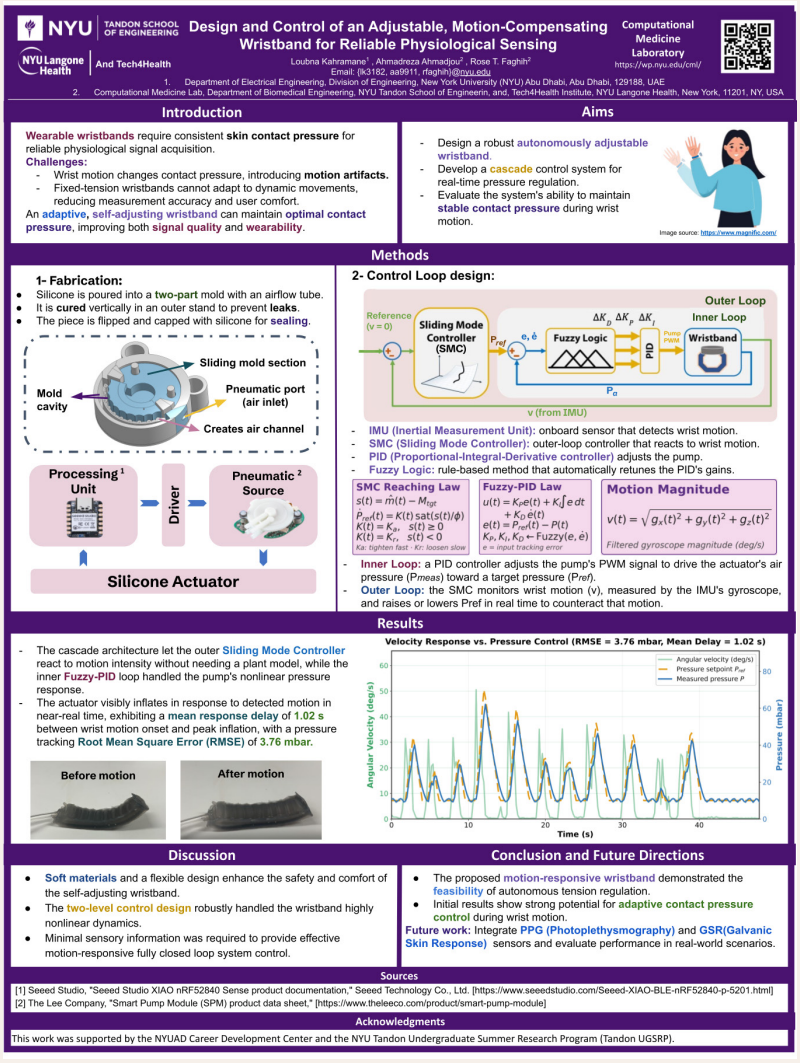


Loubna Kahramane
Electrical Engineering, B.S.
Spring 2027
NYU Abu Dhabi

Design and Control of an Adjustable, Motion-Compensating Wristband for Reliable Physiological Sensing

Background & Objective: Smart wearables capture and process physiological data to provide valuable insights about health conditions to the user. Besides environmental noise, movements from intense physical activity can cause readings that do not reflect correct information. Sensor measurement errors resulting from movement are referred to as motion artifacts. Motion artifacts can potentially overshadow the small underlying physiological signal. To mitigate the effect of motion artifacts, we propose a mechanism for smartwatches through which the wearable reacts to the motion in real time and tightens the smartwatch's wristband accordingly, which potentially leads to more reliable readings. This can also significantly simplify data preprocessing in subsequent stages.

Methods: To achieve the aforementioned, we design a wristband that changes its shape according to accelerometer data. More specifically, we incorporate a molding procedure to fabricate an actuator with a blocked air channel using silicon. Once the inner channel of the actuator is exposed to sufficient air pressure, it changes its shape by expanding and causes the wristband to get tighter around the user's wrist. To be able to achieve the required air pressure, we use a micro air pump which is controlled using a two-stage controller. More concretely, the controller features two loops: the inner loop incorporates a Fuzzy-PID controller to control the measured pressure based on the pump's driver's voltage and the outer loop is a model-free sliding mode controller which receives velocity feedback to set a time-varying baseline for the pressure.



Results: On the bench, the full cascade firmware worked as intended. When wrist motion increased, the outer loop raised the pressure setpoint, and the inner Fuzzy-PID loop tracked that setpoint by adjusting pump PWM in response, tightening the band accordingly. The final mold iteration gave us a thinner, more compact actuator design, and by the end we had a complete physical prototype: the actuator, pump, custom enclosure, and straps all assembled into one wearable band.

Conclusion: In conclusion, we have demonstrated the feasibility and potential of our proposed prototype, which integrates advanced control strategies with a purpose-driven mechanical design. Our observations show great promise for producing motion-responsive tightening behavior. The next stage of development will focus on incorporating photoplethysmography (PPG) and galvanic skin response (GSR) sensors into the existing platform and evaluating the applicability of the device in real-life scenarios to assess its responsiveness, robustness, reliability, and overall performance during different activities and environmental conditions.



Luhan Zhao
Chemical and Biomolecular Engineering, B.S
Spring 2027
NYU Tandon
AHA Fellow

Characterizing GPR83 Signaling in Regulatory T Cells in the Context of Atherosclerotic Inflammation

Atherosclerosis is a chronic inflammatory disease of the arteries, a lipid-driven process defined by the accumulation of low-density lipoprotein (LDL) in the vessel wall. Inflammatory macrophages play a large role in atherosclerotic plaque formation by entering artery walls and engulfing excess cholesterol, becoming "foam cells." Regulatory T cells, Tregs, are a specialized group of immune cells that suppress inflammation by releasing immunosuppressive cytokines into the atherosclerotic milieu, thus controlling the inflammatory macrophages and reducing plaque formation. GPR83, a recently deorphanized G-protein coupled receptor, is highly expressed by Tregs, whose function on these cells are still unclear. Herein we explore the function of this receptor via detection of a downstream mediator of receptor signaling. GPR83 signals through its partner G-protein, Gq/11, which activates the phospholipase C (PLC) signaling cascade, culminating in IP-1 accumulation. To characterize this accumulation, a known activator of the PLC cascade in T cells, concanavalin A, was used with Homogeneous Time-Resolved Fluorescence (HTRF) technology. The assay had two points of adjustment, activation duration and cell density. The results of the optimization assay with bulk CD4+ T cells showed that longer activation time and increased cell count demonstrated a greater dynamic range of signal. The next iteration of the assay included stimulation of Tregs, activated T cells, and naive T cells with a GPR83 ligand, and showed inconclusive results requiring further assay modifications to confirm GPR83 signaling. Next steps for this project include further assay optimization, confirming GPR83 receptor expression, and RNA sequencing of the Treg transcriptomic profile after GPR83 activation. Once GPR83 signaling is confirmed and the downstream effects on Treg function are characterized, we would like to employ this signaling cascade for use in therapeutic platforms.

Investigating the Role of Regulatory T Cell GPR83 Signaling in Atherosclerotic Inflammation

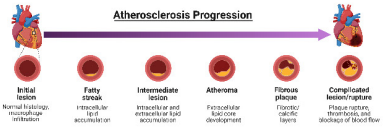
Luhan Zhao[1], Aaron T. Alpar[1], Jeffrey A. Hubbell[1]
[1]Chemical and Biomolecular Engineering, NYU Tandon

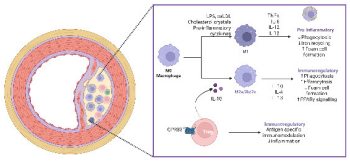
NYU TANDON SCHOOL OF ENGINEERING

SUMMARY

Atherosclerosis is a chronic inflammatory disease of the arteries, a lipid-driven process defined by the accumulation of low-density lipoprotein (LDL) in the vessel wall. Inflammatory macrophages play a large role in atherosclerotic plaque formation by entering artery walls and engulfing excess cholesterol, becoming "foam cells." Regulatory T cells, Tregs, are a specialized group of immune cells that suppress inflammation by releasing immunosuppressive cytokines into the atherosclerotic milieu, thus controlling the inflammatory macrophages and reducing plaque formation. GPR83, a recently deorphanized G-protein coupled receptor, is highly expressed by Tregs, whose function on these cells are still unclear. Herein we explore the function of this receptor via detection of a downstream mediator of receptor signaling. It is strongly suggested that GPR83 signals through its partner G-protein, $G_q/11$, which activates the phospholipase C (PLC) signaling cascade, culminating in IP-1 accumulation. To characterize this accumulation, a known activator of the PLC cascade in T cells, concanavalin A, was used with Homogeneous Time-Resolved Fluorescence (HTRF) technology. The assay had two points of adjustment, activation duration and cell density. The results of the optimization assay with bulk CD4+ T cells showed that longer activation time and increased cell count demonstrated a greater dynamic range of signal. The next iteration of the assay included stimulation of Tregs, activated T cells, and naive T cells with a GPR83 ligand, and showed inconclusive results requiring further assay modifications to confirm GPR83 signaling. Next steps for this project include further assay optimization, confirming GPR83 receptor expression, and RNA sequencing of the Treg transcriptomic profile after GPR83 activation. Once GPR83 signaling is confirmed and the downstream effects on Treg function are characterized, we would like to employ this signaling cascade for use in therapeutic platforms.

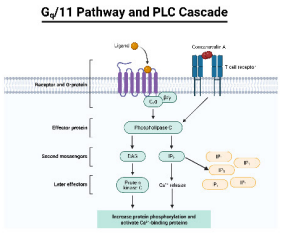
BACKGROUND





As macrophages engulf excess cholesterol and oxidized low-density lipoprotein, they become foam cells that drive atherosclerotic inflammation and arterial plaque formation[1]. Conversely, regulatory T cells, Tregs, can suppress vascular inflammation through immunoregulatory mechanisms, including the secretion of IL-10. Exogenous IL-10 has demonstrated success in reducing plaque severity[2]. GPR83, a recently deorphanized GPCR, is highly expressed by Tregs, and controlling signaling through this receptor may provide a means of modulating Treg immunoregulatory activity, potentially including endogenous IL-10 production. Therefore, we seek to establish an assay for investigating GPR83 signaling as a first step toward determining whether this pathway could be harnessed to reduce atherosclerotic inflammation.

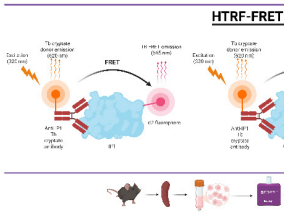
METHODS



GPR83 signals through the Gq/11 and PLC cascade[3], which generates the secondary messenger, IP3. IP3 is degraded into IP1, and its accumulation can be measured with the revvity HTRF IP-One assay.

The assay works by measuring FRET intensity between the a fluorescently labeled antibody (donor) and a labeled IP1 (acceptor), both provided by the assay. Cell derived IP1 outcompetes the labeled IP1, resulting in a reduction in FRET signal. Thus an inverse relationship between fluorescent intensity and cell derived IP1 concentration can be measured[4].

There are two aspects of the assay which can be optimized to increase the dynamic range: activation time and cell-plating density.



Spleens were isolated from C57BL/6 mice, before generating single cell suspensions. CD4+ T cells, bulk or naive, were then separated using StemCell magnetic sorting kit, before plating in differentiation or activation media, according to manufacturer instructions.

RESULTS

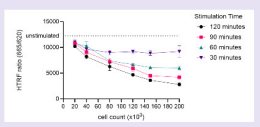


Figure 1: Bulk CD4+ T cells were isolated via magnetic separation and stimulated at the indicated time points with with the 5 ug/mL Concanavalin A. Each point represents mean of n=3 with error bars for SD.

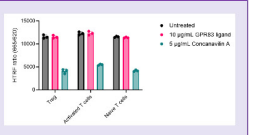


Figure 2: T cells were either cultured for 6 days (Treg and activated), or isolated day of, followed by a two hour rest. 80,000 cells were plated per condition, and stimulated for two hours. Each point represents one sample with error bars for SD.

CONCLUSION

Depicted in **Fig. 1** are the results from the optimization assay. As we increase both the stimulation time and cell density, a greater drop in fluorescence observed, denoting an increase of PLC activity. **Fig. 2** depicts the results from the GPR83 activity assay, in which no discernable differences in IP1 accumulation between untreated and GPR83-ligand induced signaling cells are observed in contrast to the cells that received the positive control, ConA.

Based on our results, it seems we did not generate GPR83 activation. We are planning on further optimizing the GPR83 signaling assay, as well as performing qPCR on our Tregs to confirm GPR83 expression. In the instance that we don't see GPR83 expression in our in vitro generated Treg, we may move forward with isolating fresh Tregs rather than culturing.

REFERENCES

(1) Wolf D, Ley K. Immunity and inflammation in atherosclerosis. Circ Res. 2019;124(2):315–327. doi:10.1161/CIRCRESAHA.118.313591.

(2) Volpatti LR, Norton de Matos S, Borjas G, Beckman TN, Reda JW, Watkins EA, et al. LDL-binding IL-10 reduces vascular inflammation in atherosclerotic mice. Nat Biomed Eng. Published online 2025. doi:10.1038/s41551-025-01573-8.

(3) Müller A, Leinweber B, Fischer J, et al. The extracellular N-terminal domain of G-protein coupled receptor 83 regulates signaling properties and is an intramolecular inverse agonist. BMC Res Notes. 2014;7:913. doi:10.1186/1756-0500-7-913.

(4) Revvity. HTRF IP-One Gq Detection Kit: Instruction Manual. Version 11. October 2025. Part nos. 62IPAPEB, 62IPAPEC, and 62IPAPEJ.

ACKNOWLEDGEMENTS

American Heart Association.

I would like to thank the NYU Tandon Undergraduate Summer Research Program and the American Heart Association for the opportunity to pursue research this summer.

Created with BioRender Poster Builder



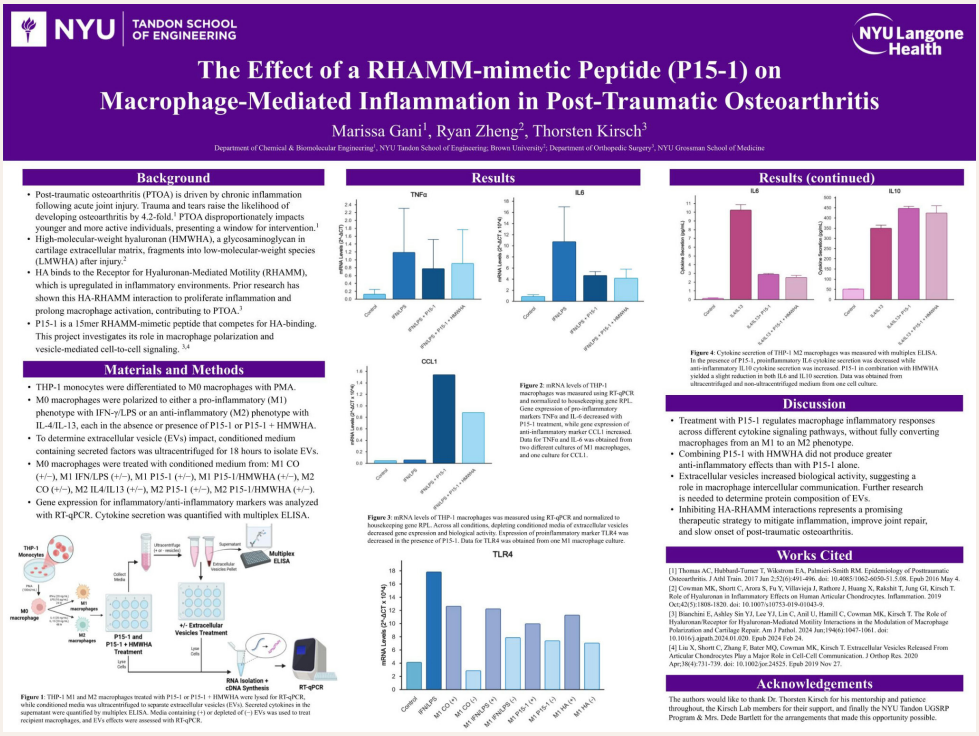
Marissa Gani
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon
Thompson-Bartlett Fellow



Ryan Zheng
Computer Science, B.S.
Spring 2029
Brown University

The Effect of a RHAMM-mimetic Peptide (P15-1) on Macrophage-Mediated

Inflammation in Post-Traumatic Osteoarthritis Persistent inflammation following acute joint injury contributes to post-traumatic osteoarthritis (PTOA), a degenerative joint disease characterized by pain, swelling, and cartilage deterioration. Hyaluronan (HA), a lubricating glycosaminoglycan in joint synovial fluid, breaks down into low-molecular-weight fragments (LMWHA) that interact with the Receptor for Hyaluronan-Mediated Motility (RHAMM), which is markedly upregulated after joint injury. Previous studies from our laboratory and others have shown that HA-RHAMM signaling promotes inflammatory immune responses in macrophages, driving progression into PTOA. Our laboratory is investigating P15-1, a 15mer RHAMM-mimetic peptide, as a potential therapeutic to disrupt vicious HA-RHAMM interactions and reduce prolonged macrophage activation. Human THP-1 monocytes were differentiated and polarized into either pro-inflammatory (M1) or anti-inflammatory (M2) macrophage phenotypes, then treated with P15-1, or P15-1 in combination with high-molecular-weight HA (HMWHA). Expression of inflammatory and anti-inflammatory genes was analyzed by RT-qPCR, and cytokine secretion was quantified by multiplex ELISA. In addition, conditioned medium from M1 macrophages was used to investigate macrophage crosstalk through extracellular vesicles. Stimulating macrophages with interferon-gamma (IFN˜) and lipopolysaccharide (LPS), as well as conditioned medium containing cell-secreted factors from M1 macrophages, induced an inflammatory phenotype. Treatment with P15-1 reduced expression of several inflammatory markers and decreased the release of pro-inflammatory cytokines, but P15-1 in combination with HMWHA did not produce greater anti-inflammatory effects than P15-1 alone. Removing macrophage extracellular vesicles from conditioned medium decreased biological activity, suggesting that extracellular vesicles play a role in intercellular communication. These findings demonstrate that P15-1 inhibits HA-RHAMM interactions and modulates macrophage inflammatory responses, without fully converting from an M1 to M2 phenotype. Overall, our research provides insights into a new promising therapeutic strategy to control the inflammatory microenvironment after joint injury, ultimately improving tissue repair and slowing the progression of PTOA.





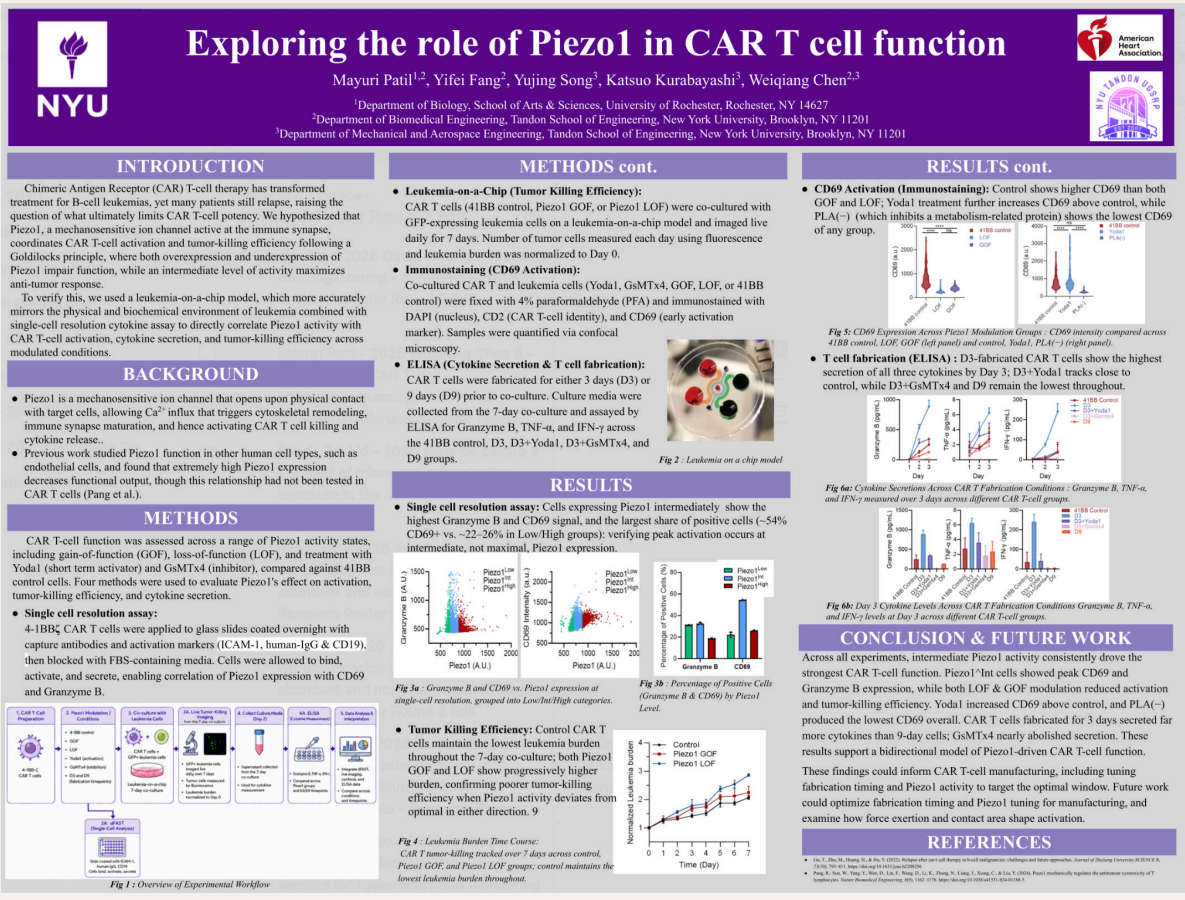
Mayuri Patil
Biomolecular Science, B.S.
Spring 2029
University of Rochester
AHA Fellow

Exploring the role of Piezo1 in CAR T cell function motivation

Piezo1 is a mechanosensitive ion channel that regulates T-cell function by coupling cytoskeletal dynamics to metabolic rewiring. Previous studies have shown that Piezo1 disruption impairs glycolytic and mitochondrial ATP production, leading to energetic stress and reduced cell cytotoxicity. These findings suggest that Piezo1 activity follows a "Goldilocks" principle in which an intermediate level of activity maximizes T-cell activation and cytotoxic function, whereas insufficient or excessive activity compromises T-cell fitness. This study investigates the role of Piezo1 in regulating CAR T-cell killing efficiency using a leukemia-on-a-chip model.

CAR T cells with Piezo1-activated, Piezo1-inhibited, gain-of-function (GOF), and loss-of-function (LOF) conditions were evaluated using a leukemia-on-a-chip platform. Cell culture media were collected from the chip model over seven days. The concentrations of Granzyme B, IFN-γ, and TNF-α were quantified using ELISA to evaluate the cytotoxic capacity and tumor-killing efficiency of CAR T cells. Fixed chip devices were immunostained with DAPI, CD2, and CD69 and imaged using confocal laser scanning microscopy (CLSM). CD69 fluorescence intensity on CD2+ cells was quantified across Piezo1 groups. CAR T cells were also cultured on CD19/ICAM-coated micropillar arrays that mimic the CD19+ tumor cell surface and image analysis quantified the mechanical forces exerted by activated CAR T cells.

ELISA demonstrated that cytokine secretion was positively associated with the mechanotransduction capacity of CAR T cells across Piezo1 groups. Micropillar analysis showed that Piezo1 activation increased the mechanical forces exerted by CAR T cells, suggesting enhanced immunological synapse formation that may lead to improved tumor cell-killing efficiency. Compared with the Piezo1-inhibited and LOF groups, the Piezo1-activated and GOF groups exhibited greater cytokine secretion, increased traction force generation, and enhanced tumor cell-killing efficiency. These findings suggest that Piezo1 modulation enhances CAR T-cell mechanotransduction and may contribute to future CAR T-cell sorting or engineering strategies to achieve improved therapeutic outcomes.



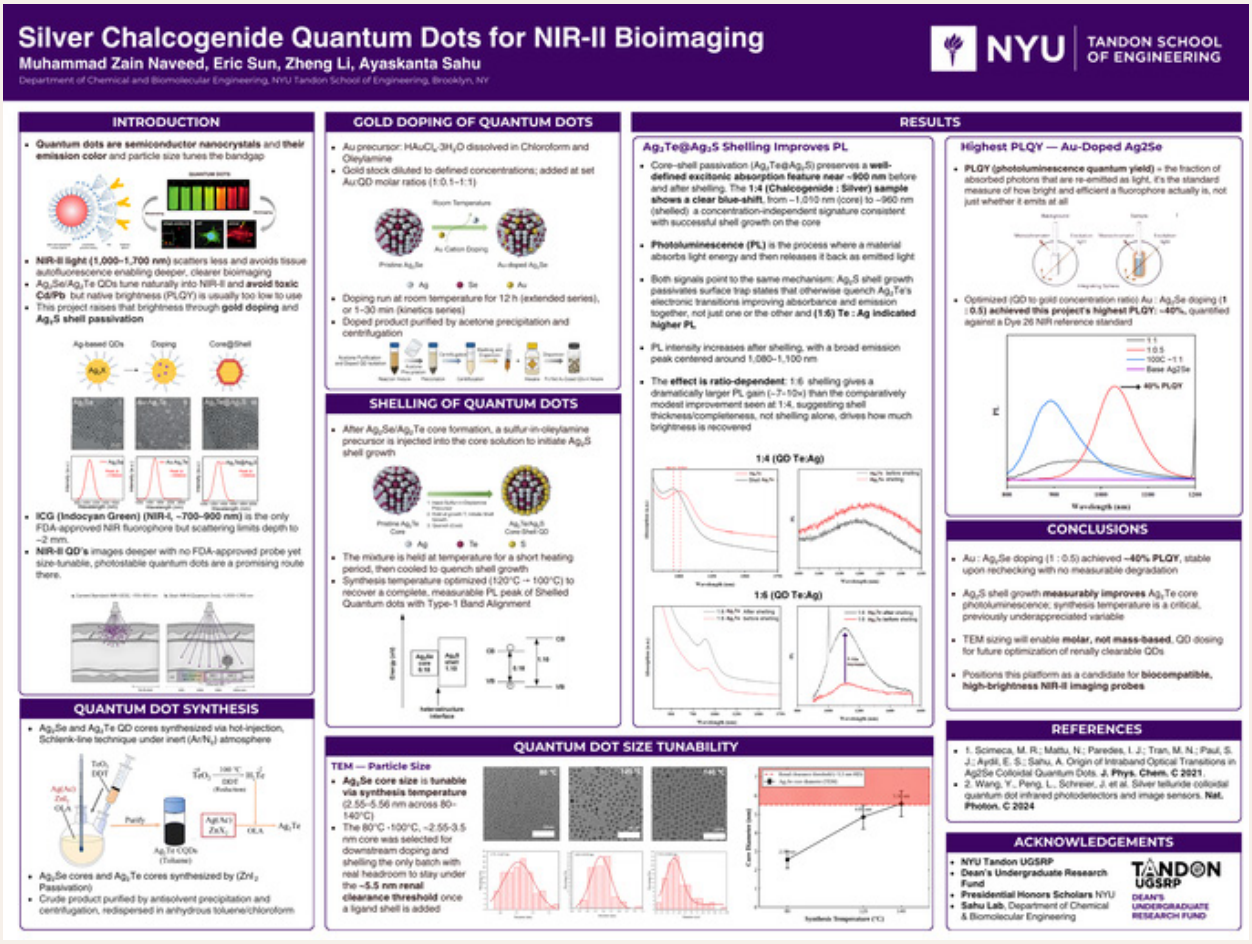
Engineering Health



Muhammad Zain Naveed
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU CAS

Core-Shell Passivation and Gold Doping of Silver Chalcogenides Quantum Dots for NIR-II Cancer Diagnostic Imaging

Fluorescent probes that emit in the second near-infrared window (NIR-II, 1,000–1,700 nm) enable non-invasive, deep-tissue cancer imaging with reduced tissue scattering and autofluorescence, yet clinically viable NIR-II emitters with a high photoluminescence and biocompatibility remain scarce. Here we engineer silver chalcogenide quantum dots (Ag_2Se, Ag_2Te), a low-toxicity alternative to conventional heavy-metal quantum dots, through core–shell passivation and gold doping to substantially increase their near-infrared emission efficiency, working toward biocompatible, PEGylated probes for future cancer-targeted imaging. Quantum dot cores were synthesized via hot-injection methods under Schlenk-line conditions, followed by Ag_2S shell growth to passivate surface trap states and gold doping (Au:Ag_2Se) to introduce mid-gap states that enhance near-infrared emission. Core–shell passivation of Ag_2Te with Ag_2S produced a marked increase in photoluminescence intensity relative to the unshelled core, confirming that shell growth suppresses non-radiative surface trap states. Separately, an optimized gold-doping route (Au:Ag_2Se, 0.5:1) achieved a PLQY of approximately 40%, with no measurable degradation upon rechecking. Subsequent efforts to systematically replicate this doping result across a wider range of ratios, carrier solvents, and reaction temperatures did not reproduce comparable PLQY, indicating that quantum dot concentration determination and precursor delivery conditions are critical, underappreciated variables in this synthesis. These findings establish two complementary strategies, shell passivation and gold doping, for improving NIR-II quantum dot brightness, and inform ongoing efforts to identify a reproducible, high-PLQY batch suitable for future PEGylation and biocompatibility studies. Together, this work advances the synthetic groundwork for translating silver chalcogenide quantum dots into clinically relevant, high-brightness NIR-II probes for early cancer detection.



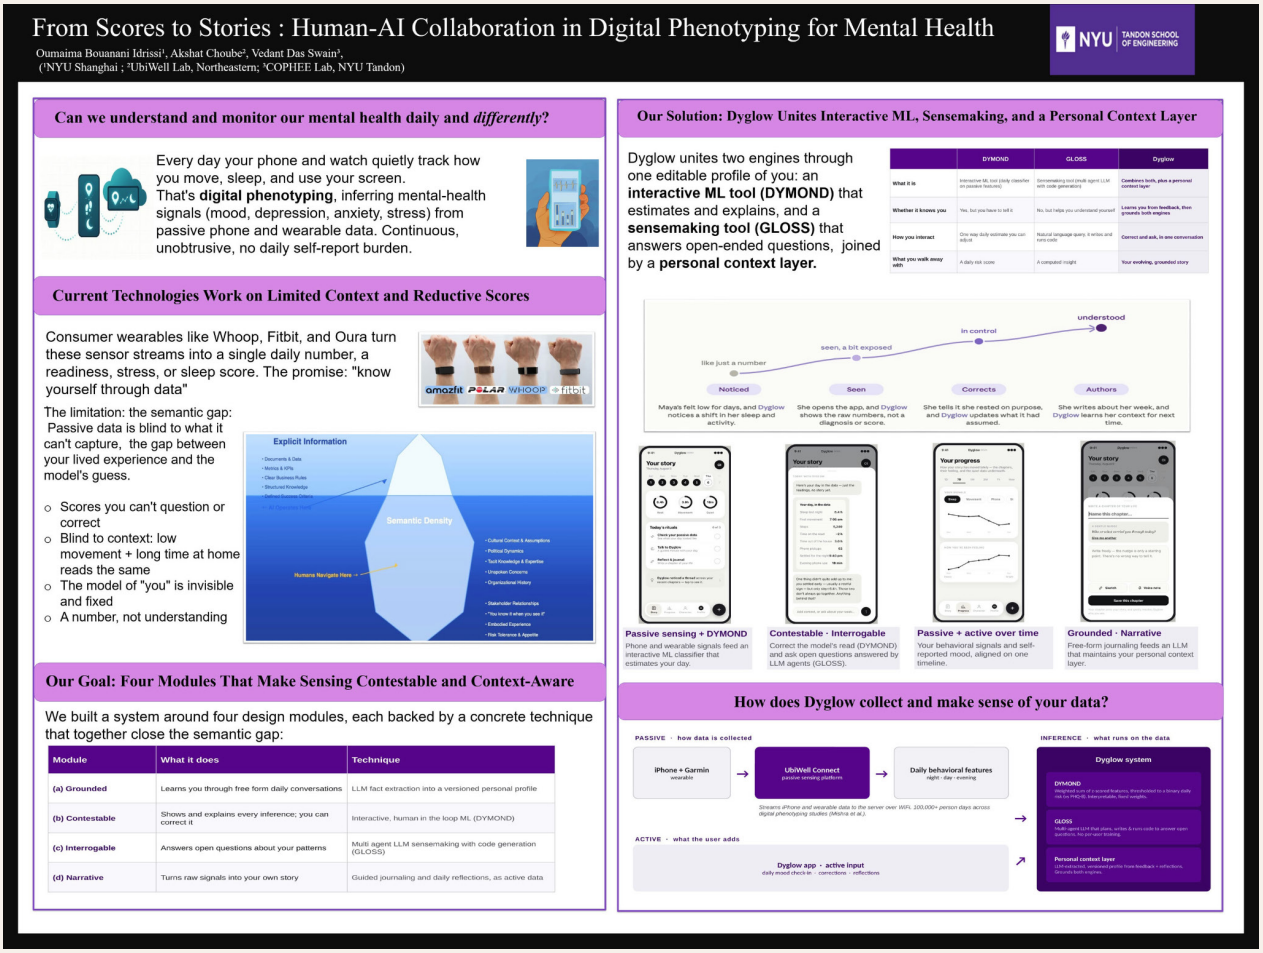
Engineering Health



Oumaima Bouanani Idrissi
Data Science, B.S.
Spring 2027
NYU Shanghai

An AI Companion that Turns Passive Mental-Health Sensing into a Correctable, Narrative Practice

Smartphones and wearables have become adept at passively inferring behavioral indicators like physical activity, screen use, and sleep, yet translating these signals into reliable models of mental health outcomes remains an elusive challenge. And where such systems do attempt it, they typically treat the person as a subject to be measured: their inferences are invisible, cannot be corrected, and ignore the lived context that gives sensor patterns meaning, offering little to help users actually understand their own wellbeing. This project develops a mobile system, DYGLOW, whose goal is to help people better understand their mental health by reframing daily depression-risk sensing from a measurement problem to a collaborative practice mediated by an AI companion. It builds on two existing engines for behavior monitoring, (i) a classifier-based model that produces a daily risk-oriented insight from passive features, and (ii) an agentic large-language-model system that performs open-ended analysis of the same data. DYGLOW integrates these engines through a shared "personal model" that stores user feedback (corrections, life context, and preferences) as a mutable natural-language profile. This model grounds both engines, so the system's representation of each user becomes progressively more personalized while remaining transparent and user-controlled. Through a single conversational interface, users view their raw signals, resolve inconsistencies the model surfaces, report how each day actually felt, ask questions about their own patterns, and compose guided daily "chapters" via text, sketch, and voice, every interaction both teaching the model and helping the user reflect on and make sense of their own mental health. We implemented the full system and a working prototype driven by test data, verifying the feedback-to-inference loop end to end. Our approach suggests that passive-sensing systems can help people become active, informed participants in understanding their own mental health, rather than passive subjects of measurement.





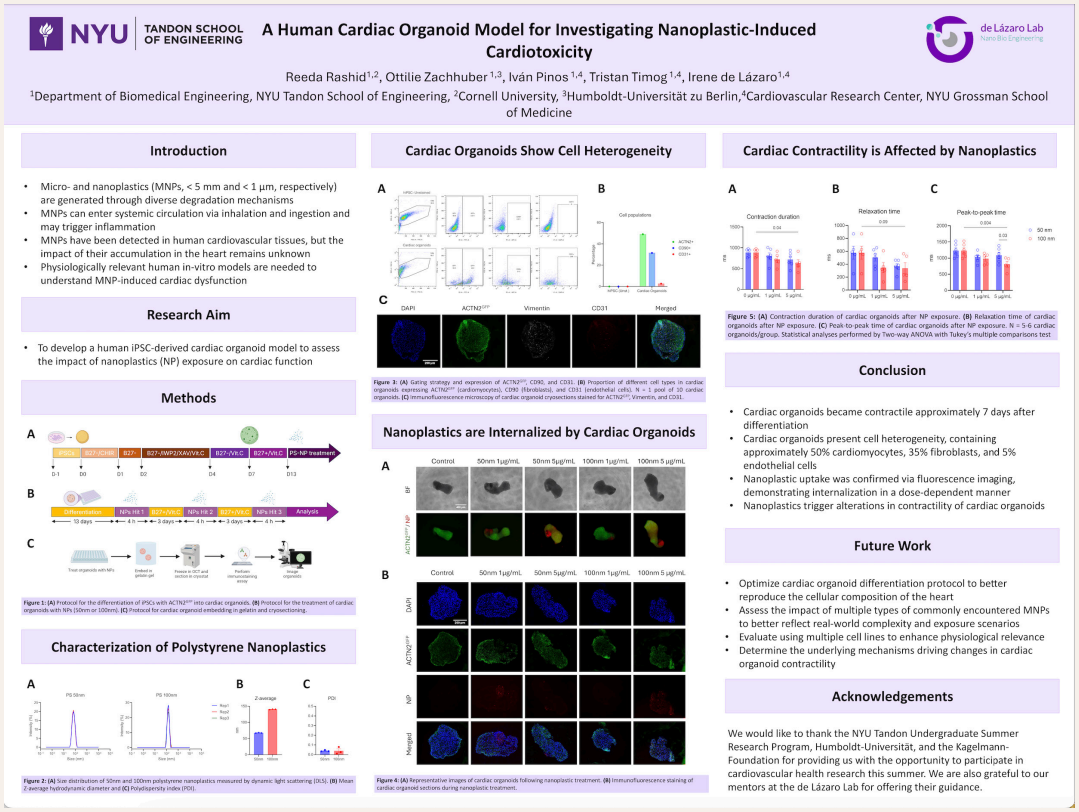
Reeda Rashid
Biological Sciences, B.S.
Spring 2028
Cornell University



Ottilie Zachhuber
Biology, B.Sc
Winter 2026
Humboldt University

A Human Cardiac Organoid Model for Investigating Nanoplastic-Induced Cardiotoxicity

Plastic pollutants are now recognized as ubiquitous environmental hazards, with micro- and nanoplastics (MNPs, < 5 mm and < 1 µm, respectively), generated through the degradation of larger plastic waste by weathering, UV radiation, and mechanical abrasion. Recent evidence has detected MNPs in human tissues and the bloodstream, suggesting these particles can enter systemic circulation via inhalation and ingestion and trigger inflammation. These findings have raised concerns about the potential role of MNPs as a driver of cardiovascular disease. However, our understanding of the mechanisms underlying MNP-induced cardiac dysfunction is limited by the lack of physiologically relevant human in-vitro models. This study aims to address this research gap by developing a cardiac organoid model to assess the impact of NPs in cardiac function. We optimized a protocol to differentiate iPSC-derived cardiac organoids and characterized them using flow cytometry and immunohistochemistry analyses. To model chronic NP exposure in the heart, cardiac organoids were repeatedly exposed (4 hours every 3 days for a total of 9 days) to two different sizes of fluorescently labeled polystyrene nanoparticles, 50 nm and 100 nm in diameter. NPs uptake and distribution within the organoids were assessed using live microscopy and immunofluorescence imaging. Finally, contractile function was evaluated using MuscleMotion analysis. Cardiac organoids were functionally validated by contraction approximately 7 days after differentiation and demonstrated a composition of approximately 50% cardiomyocytes, 35% fibroblasts, and 5% endothelial cells. Nanoplastic uptake was confirmed by live-cell microscopy and cryosection imaging, with the smaller NPs (50 nm) being more abundantly internalized by the organoids. Nanoplastic uptake and distribution within cardiac organoids was confirmed by live-cell microscopy and cryosection imaging. Altogether, our results show that cardiac organoids are functional and recapitulate the cellular composition of developing human hearts. In addition, NPs are taken up by these organoids, and future studies will assess their potential detrimental effects.





Sulaiman Rasuli
Mechanical Engineering, B.S.
Spring 2030
NYU Tandon

Design and Validation of a Toggle Clamp System for Reliable Sealing of a Microfluidic Point-of-Care Diagnostic Cartridge

Patients that are critically ill often require point-of-care diagnostic platforms that must provide intuitive, user-friendly operations to enable efficient use in clinical settings. For platforms utilizing disposable microfluidic cartridges, reliable operation depends on precise control of fluid transport, which in turn requires robust fluidic sealing between the cartridge and instrument manifold. Consequently, the cartridge engagement mechanism must simultaneously deliver uniform compressive force across the cartridge interface to maintain seal integrity while remaining simple and ergonomic for the user. The aim of this research project focuses on upgrading the sealing assembly the automated cartridge-based platform from a screw-based compression system to proper sealing with a custom 3D-printed clamping plate to improve both sealing performance and accessibility. These tests require adequate pressure to the cartridge to ensure that a microfluidic seal is generated for fluid to go through the cartridge. To evaluate the validity of the design, static equilibrium and linkage tests were used to model the relationship between the applied force and the resulting compressive seal across the cartridge's interface. The design was verified using fluid dye tests that represent how fluid would move during an actual sample test, which prevented fluid leakage during testing and improved the ergonomics of the cartridge engagement. The new seal demonstrated an effective balance between mechanical performance and intuitive platform operation.

NYU TANDON SCHOOL OF ENGINEERING

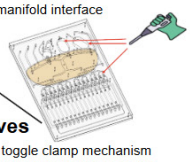
Design and Validation of a Toggle Clamp System for Reliable Sealing of a Microfluidic Point-of-Care Diagnostic Cartridge

[1]Department of Mechanical and Aerospace Engineering, and [2]Department of Biomedical Engineering

Sulaiman Rasuli[1], Michael Forte[1], Adrienne D. Füredi[2], and Katsuo Kurabayashi[1]

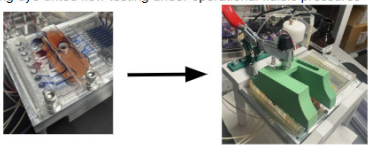
Introduction

- Critical care medicine requires near-patient diagnostics with low-complexity workflows to minimize time-to-result and operator variability [1]
- Precise volumetric fluid transport relies on a pressure-resilient, leak-tight seal across the cartridge-manifold interface
- Screw-driven clamping mechanism introduces operator error and delay in setup time



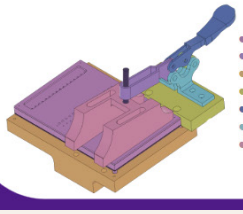
Objectives

- Develop an ergonomic, over-center toggle clamp mechanism featuring a custom force-distributing clamping plate to deliver uniform compressive force across the cartridge interface
- Perform static equilibrium and kinematic linkage analyses to quantify the mechanical advantage and establish the relationship between operator input force and output sealing force
- Characterize seal integrity, channel routing, and fluid containment using dye-tinted flow testing under operational fluidic pressures



Assembly

- Custom 3D-printed mounting and pressure plates were retrofitted onto a standard toggle clamp and rigidly fastened to the cartridge manifold



Methods

- **Force & Seal Modeling:**
  - Modeled the relationship between applied force and compressive seal integrity using static equilibrium and linkage testing
- **Pressure & System Comparison:**
  - Evaluated clamping force and pressure gradients using a 55 x 85 mm rubber spring (stiffness k = 92.1 N/mm)
  - Compared performance metrics against the previous system to quantify improvements.

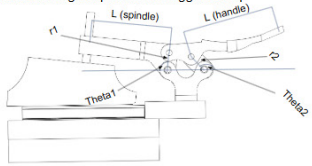
F_c = Clamping Force, Δx = amount of compression, k_{sys} = Combined Stiffness

$$F_c = -k\Delta x$$

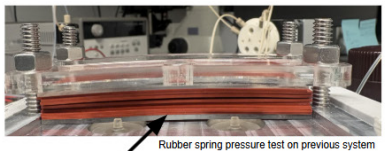
$$P = \frac{F_c}{A} \mid A = Surface\ Area,\ P = Pressure$$

$$F_c = F_i\left(\frac{L_{handle}}{L_{spindle}}\right)\left(\frac{r_2 \sin\theta_2}{r_1 \sin\theta_1}\right)$$

- Mechanical Advantage Equation for Toggle Clamp



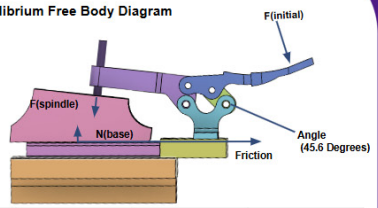
- **Clamping Force & Pressure Analysis:**
  - Estimated clamping force via Hooke's Law using 3-point (left, center, right) compression data to guide design iterations.
  - Calculated cartridge pressure using known contact surface area.
- **Design Verification:**
  - Conducted dye testing under active fluid flow to verify microfluidic seal integrity.



- Previous system providing uneven pressure, making the center bulge

Results

Static Equilibrium Free Body Diagram



	Screw-based System	Toggle-Clamp System Ver 3.1	Toggle-Clamp System Ver 3.2
F on CENTER	0 N (Bulging upwards)	45.90 N	35.92 N
F on RIGHT SIDE	35.92 N	79.82 N	99.78 N
F on LEFT SIDE	35.94 N	11.97 N	23.95 N
Net Force	71.84 N	137.69 N	159.64 N

- **Pressure Distribution Analysis:**
  - Previous System exhibited a central pressure deficit when tested with the rubber spring
  - Redesigned System displayed an asymmetric pressure gradient, with minimum pressure localized on the left side
- **Failure Mode Verification:**
  - Dye Testing confirmed that the left-side pressure deficit compromised manifold seal integrity, inducing sample leakage during loading



Leakage on the left connection

Conclusion and Future Work

- Rotational toggle actuation creates an **asymmetrical friction profile**, driving premature left-side engagement. This creates a localized pressure deficit that leads to manifold leakage during fluid loading.
- Transition from a rotational mechanism to a pure **vertical linear actuation** system to eliminate friction asymmetry and ensure simultaneous, uniform clamping across the cartridge.
- To validate it, pressure distribution modeling and active dye leakage testing will be performed on the linear assembly to confirm microfluidic seal integrity.

References

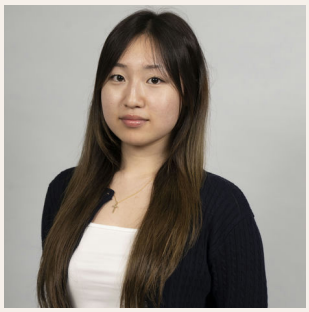
[1] Füredi, A.D. *et al.* An automated, digital immunoassay on a microfluidic cartridge for on-demand cytokine profiling. *Microsyst Nanoeng* **12**, 252 (2026).

August 7th 2026, UGSRP Research Poster Symposium

Engineering Health



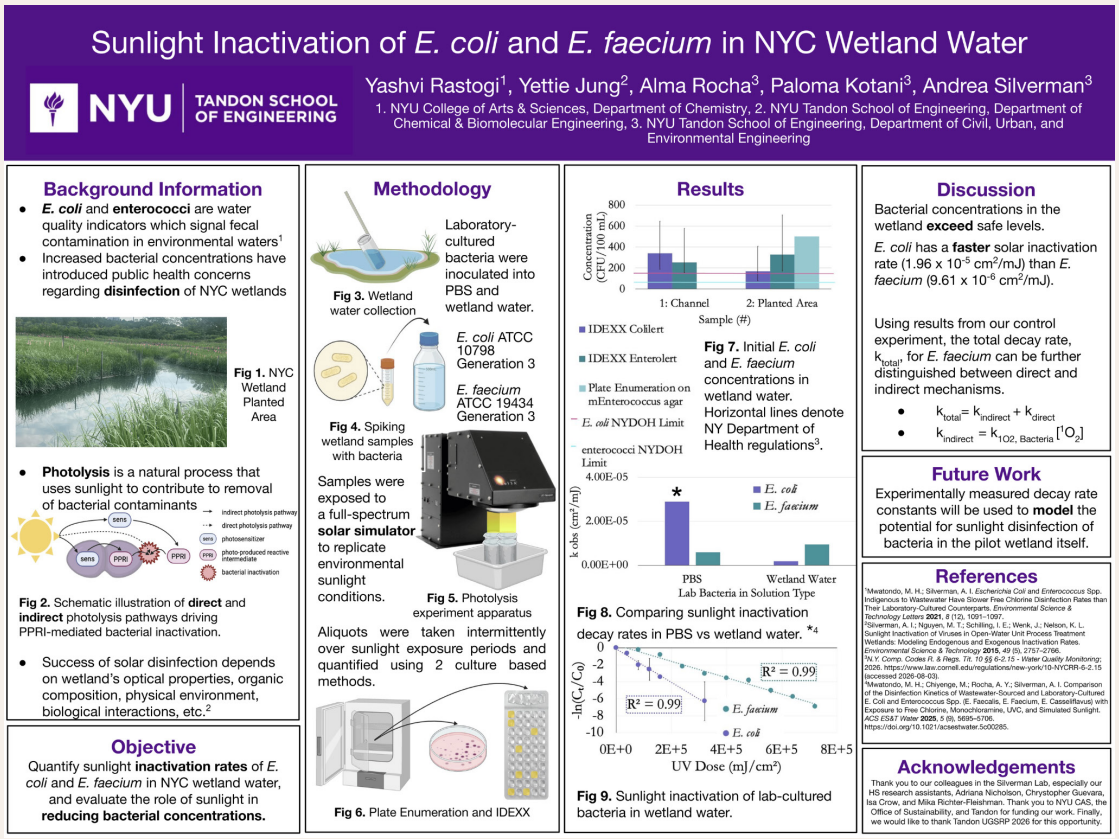
Yashvi Rastogi
Chemistry, B.S.
Fall 2027
NYU CAS



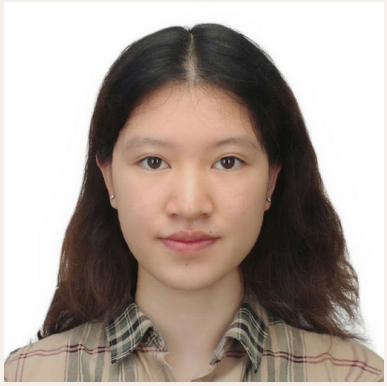
Yettie Jung
Biomolecular Science, B.S.
Spring 2028
NYU Tandon
Thompson-Bartlett Fellow

Sunlight Inactivation of Enterococcus and E. coli in NYC Wetland Water

Enterococcus spp. and Escherichia coli are commonly used as indicators of fecal contamination in environmental waters. Increased concentrations of these bacterial contaminants in NYC wetlands have raised concerns regarding public health and current disinfection practices. Sunlight inactivation, also known as photolysis, is a natural process that contributes to the removal of bacterial contaminants through both direct and indirect photolysis mechanisms, utilizing photo-produced reactive intermediates and photosensitizers. However, the success of disinfection heavily depends on environmental conditions, including water optical properties and the presence of organic compounds such as dissolved organic matter. This project aimed to quantify the sunlight inactivation rates of Enterococcus spp. and E. coli in NYC wetland water and evaluate the role of sunlight in reducing these bacterial concentrations. Wetland water was sampled and spiked with reference strain E. coli (ATCC 10798) and E. faecium (ATCC 19434), and exposed to a full-spectrum sunlight solar simulator to replicate environmental sunlight conditions. Bacteria were expected to undergo both direct and indirect photolysis mechanisms. Dark controls were unexposed to light to account for potential non-photolytic bacterial decay. Samples were collected periodically over the course of the experiment and quantified using 2 culture-based methods: plate enumeration on mTEC and mEnterococcus agar, and IDEXX Colilert and Enterolert (51-Well Quanti-Tray). Experimental results show that E. coli has a faster solar inactivation rate than E. faecium, with rate constants of 1.96x10-5 cm2/mJ and 9.61x10-6 cm2/mJ, respectively. Results from a control experiment were used to distinguish the total inactivation rate for E. faecium, ktotal, into direct and indirect photolysis components. The direct and indirect inactivation rates were 5.96 x 10-6 cm2/mJ and 3.65 x 10-6 cm2/m respectively, summing to its total inactivation rate. In the future, these experimentally measured decay rate constants will be used to model the potential for sunlight disinfection of bacteria in the pilot wetland itself.



Engineering Health



Yichen Zhao
Mathematics, B.S.
Spring 2027
NYU Shanghai

Towards Robust Long-Term Neural Decoding: Baselines and Benchmark on the LINK Dataset

Long-term stability remains one of the primary challenges in intracortical brain–computer interface (BCI) systems due to neural signal variability across recording sessions. In this project, we establish a benchmarking pipeline for evaluating long-term neural decoding on the LINK dataset, which contains neural recordings collected over hundreds of days. We implement and evaluate several representative approaches, including NoMAD, StabilizedBCI (Degenhart et al.), ERDiff, SPINT, and Flow Domain Alignment (FDA), under a unified evaluation framework. Their decoding performance and robustness are compared using consistent metrics to better understand the strengths and limitations of current methods for long-term neural decoding. This benchmark provides a foundation for future research on developing more robust and generalizable neural decoding algorithms.

NYU TANDON SCHOOL OF ENGINEERING

Do Cross-Day Alignment Methods Survive Years of Neural Drift?

A Unified Multi-Year Benchmark of iBCI Decoder Adaptation

Yichen Zhao,[1] Yao Wang,[2] Tianyu He[2]

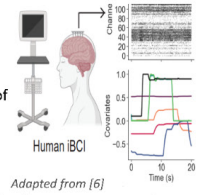
1. Undergraduate Researcher, NYU Shanghai 2. Department of Electrical and Computer Engineering, NYU Tandon School of Engineering

Abstract

We present the first unified benchmark of representative long-term neural decoding methods on the 3.5-year LINK dataset. Using identical preprocessing, decoder, and evaluation protocols, we compare multiple alignment- and architecture-based approaches under the same experimental setting. Surprisingly, none consistently outperformed a supervised LSTM baseline, highlighting the remaining challenge of robust long-term neural decoding.

Background & Motivation

Intracortical brain-computer interfaces (iBCIs) decode intended movements from motor cortical activity [5]. However, neural recordings drift over time because of changes in recorded neurons and electrode interfaces, causing decoding performance to degrade across recording sessions [1,4,7].

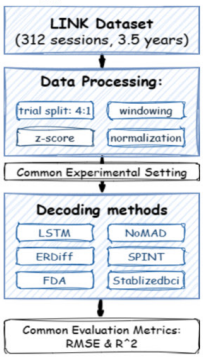


Adapted from [6]

Many methods have been proposed to improve cross-session robustness [1,2,3,4,6], but they are typically evaluated on different datasets, over different time spans, and with different decoders. As a result, reported improvements are difficult to compare directly [5].

Research question
Can existing "robust" neural decoders remain robust over years rather than weeks?

Approach



Dataset
- LINK: 312 sessions across 303 recording days (≈3.5 years)
- 2 64-channel Utah arrays
- Spiking-band power (20 ms bins), 2-DOF finger movement

Unified Protocol
- Shared preprocessing
- 13 fixed source days; target sessions within ±100 days
- 300 held-in / 75 held-out trials per session

Evaluation
- Common decoder and preprocessing
- Representative baseline and adaptation methods
- Metrics: R^2, Correlation, MSE

Figure 1. The shared pipeline. Only the decoding block changes.

Results

1. Decoding performance deteriorates over time

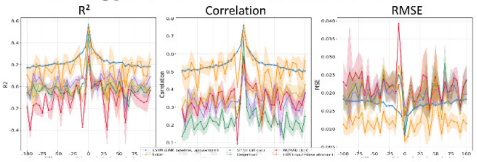


Figure 2. R^2, correlation and MSE versus the training-to-test gap.

The supervised LSTM baseline achieved $R^2 \approx 0.57$ on the source day but dropped to ≈0.22 within 25 days. None of the evaluated adaptation methods consistently exceeded the baseline across recording gaps.

Results, continued

2. Model capacity outweighs adaptation strategy

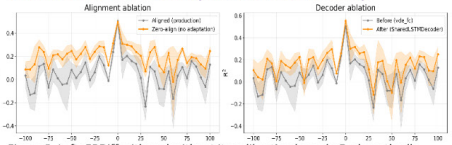


Figure 3. Left: ERDiff with and without its calibration branch. Each method's alignment stage was deleted, holding architecture and split fixed. Right: ERDiff with new decoder (versus the training-to-test gap).

Replacing shallow decoder heads with a one-layer LSTM substantially improved performance across multiple methods, whereas changing the alignment mechanism alone produced much smaller gains. This suggests that decoder capacity contributes more to performance than adaptation strategy.

Key Insights

- No published method beat a plain supervised LSTM over 3.5 years.
- Several trail the baseline at zero gap, before any drift.
- Deleting alignment did not hurt uniformly. Decoder capacity, spanning 2,500x, tracked performance instead.
- The reported gains are confounded with decoder choice and short horizons.

Future work

- Run the same protocol on other datasets
- Design an alignment method built for multi-year drift

Works Cited [1] Degenhart et al., Nat. Biomed. Eng. 4:672-685, 2020 [2] Wang, Wu, Li, Wu. ERDiff. NeurIPS 36, 2023.[3] Wang, Qi, Wang, Pan. Flow matching for few-trial neural adaptation. ICML, PMLR 267:64059-64083, 2025.[4] Karpowicz et al. NoMAD. Nat. Commun. 16:4662, 2025.[5] Karpowicz et al. FALCON Benchmark. NeurIPS 2024, Datasets and Benchmarks.[6] Le et al. SPINT. NeurIPS 2025.[7] Temmar et al. LINK Dataset. NeurIPS 2025, Datasets and Benchmarks.

Acknowledgement The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of this project, and the Chestek Lab at the University of Michigan for releasing the LINK dataset.



Yoonjae Andrew Joung
Computer Science, B.S.
Fall 2030
NYU-KAIST

Detection of Mild Cognitive Impairment via Phone-Call Network Traffic

Longer pauses and slower speech are among the earliest behavioral markers of mild cognitive impairment (MCI), but detecting them today requires dedicated audio recording in clinical settings — sparse, costly, and capturing the most privacy-sensitive data a person produces. We ask whether the same markers survive into the encrypted network traffic that voice calls already generate, a signal that is passive, always-on, and content-free. In controlled calls with synthesized speech, pause duration and speaking rate were both recoverable from uplink traffic alone (pause–downtime ρ ≈ 0.48, p < .001; 70 vs 210 WPM separable at KS ≈ 0.22, p ˜ .001), while pitch was only weakly recoverable and vocal jitter not at all — the codec preserves when speech happens and discards its spectral detail. To test this on clinical speech we turned to I-CONECT, a conversational-engagement trial in adults 75+, inventorying 39,540 recordings from 151 participants and validating 11,849 participant/staff/full-conversation triads. Extracting ground-truth speech timing required solving acoustic crosstalk between the two microphones: energy-based voice activity detection cannot separate a researcher utterance at −34.2 dB from participant bleed at −33.8 dB, so we combine pyannote speaker diarization, which removes the leaked voice by speaker identity (98.5 % recall vs 91.5 % for a single threshold), with a low-threshold rVAD pass that restores the uniform pause criterion diarization turns lack. Applied across the archive, this segments 2,168 hours of participant speech from 11,849 recordings. These recordings are then replayed through a monitored Jitsi rig to generate the paired packet traces the trial never collected, giving each clinically labeled conversation a matching traffic signal. Classification of MCI versus normal cognition across 60 balanced participants (32 MCI / 28 NC) is in progress.



NYU TANDON SCHOOL OF ENGINEERING

TANDON UGSRP

Passive Detection of Mild Cognitive Impairment via Video Call Speech Network Traffic

Yoonjae Andrew Joung[1], Rameen Mahmood[2], Gloria Xin[2], Bianca Chan[2], Bethany Clements[2], Danny Huang[3]

1. KAIST 2. NYU Tandon School of Engineering 3.PhD, NYU Tandon School of Engineering

I. Abstract

Longer pauses and slower speech are among the earliest behavioral markers of mild cognitive impairment (MCI), but detecting them today requires dedicated audio recording in clinical settings — sparse, costly, and capturing the most privacy-sensitive data a person produces. We ask whether the same markers survive into the encrypted network traffic that voice calls already generate, a signal that is passive, always-on, and content-free. In controlled calls with synthesized speech, pause duration and speaking rate were both recoverable from uplink traffic alone (pause–downtime ρ ≈ 0.48, p < .001; 70 vs 210 WPM separable at KS ≈ 0.22, p ≪ .001), while pitch was only weakly recoverable and vocal jitter not at all — the codec preserves temporal detail and discards its spectral detail. To test this on clinical speech we turned to I-CONECT, a conversational-engagement trial in adults 75+, inventorying 39,540 recordings from 151 participants and validating 11,849 participant/staff/full-conversation triads. Extracting ground-truth speech timing required solving acoustic crosstalk between the two microphones: energy-based voice activity detection cannot separate a researcher utterance at −34.2 dB from participant bleed at −33.8 dB, so we combine pyannote speaker diarization, which removes the leaked voice by speaker identity (98.5 % recall vs 91.5 % for a single threshold), with a low-threshold rVAD pass that restores the uniform pause criterion diarization turns lack. Applied across the archive this yields 352.9 hours of segmented participant speech. These recordings are then replayed through a monitored Jitsi rig to generate the paired packet traces the trial never collected, giving each clinically labeled conversation a matching traffic signal. Classification of MCI versus normal cognition across 60 balanced participants (32 MCI / 28 NC) is in progress.

II. Introduction

1. Speech

Mild Cognitive Disease (MCI) → Alzheimer's disease and related dementias (ADRD)

Speech Marker	Healthy Controls	MCI	AD
Pause length at natural breaks [1]	~805 ms / ~699 ms	—	~1296 ms / ~1215 ms
Speaking rate (phonemes/s) [2]	~7.40	~6.70	~5.92
Speaking rate (words/s) [2]	~2.51	—	~1.83
Fundamental frequency F0 (Hz) [3]	205.8 ± 29.4	—	228.6 ± 41.2
Jitter (cycle-to-cycle F0 variation, %) [3]	0.86 ± 0.68	—	1.95 ± 1.45
Shimmer (amplitude variation, %)	3.89 ± 1.97	—	5.26 ± 2.63

2. Network Traffic

Audio-based Speech Marker	Network Traffic
• Require dedicated recording equipment • Require Controlled, clinic-based settings • Storage of privacy-sensitive data	• No recording elicited after one-time installation • Can detect within-person trajectories • Preserve privacy by only recording packets • Scalable at a low cost

Of course, a packet trace is a much weaker proxy for speech than a waveform, and this work **does not claim to beat audio-based models**. The claim under test is narrower and more useful: **if enough speech-timing structure survives** into encrypted call traffic that MCI-related change is detectable from it — and **if so, the resulting biomarker is one that can be deployed continuously, at scale, in real homes.**

III. Prior Studies

1. I-CONECT data [4]

I-CONECT is a multisite Phase II RCT led by Hiroko H. Dodge testing whether frequent, cognitively stimulating conversation buffers cognitive decline in socially isolated adults aged 75+.

VC: 30-min semi-structured video conversations, WC: 10-min telephone check-in conversation

186 participants	86 (NC)	Experimental	43	**Induction** VC: 4×/week	**Maintenance** VC: 2×/week
			51	WC: 1×/week	
	100 (MCI)	Control	43		
			49	WC: 1×/week	

2. Predicting MCI from I-CONECT data [5]

Used I-CONECT audio and transcripts (Week 1 Day 3 VC) to classify MCI vs. NC.

Feature	Linguistic	Acoustic	Ensemble
AUC score	74.9	64.9	82.7

3. RouterSense: behavior from network traffic [6]

RouterSense	Network Traffic Metadata (NTM)
Platform passively captures home NTM • low-cost • device-agnostic • always-on	• packet timing • uplink/downlink byte volume • contacted domains • without touching message content

IV. Part I: Does Acoustic markers survive into network traffic

1. Design

Simulated conversations over consumer voice-call applications (WhatsApp, FaceTime)

macOS TTS → .wav Audio file with varied • Speech rate, pause structure • pitch, vocal jitter → WhatsApp, FaceTime → NTM → Temporal Activity Representation

2. Result

The traffic envelope visibly tracks the speech envelope. Temporal signal survives.

Speech parameter	Observed pattern in NTM	Quantitative evidence	?
Pause duration	Longer pauses produce extended intervals of reduced uplink activity; shorter pauses yield more continuous traffic	$p \approx 0.48$, $p < .001$, $N = 54$; ≈ 90 ms additional low-activity per 1 s of pause	✓
Speaking rate (WPM)	Faster speech → more frequent, lower-amplitude uplink bursts; Slower speech → fewer bursts with higher peaks	70 vs 210 WPM distributions separable KS ≈ 0.22, $p \ll .001$, large Wasserstein distance.	✓
Pitch / F0	Higher F0 shifts the codec's average packet size	Mean packet size 196 → 209 bytes, Cohen's d= −0.91, Mann–Whitney p = .038, ANCOVA p = .017. Bitrate weaker (d = −0.51, p = .09); all timing features non-significant	∧
Jitter	No detectable packet-level effect	Injected vocal jitter 0.86 vs 2.0 (the healthy-vs-AD range); all features $\lvert d \rvert < 0.44$, Mann–Whitney and ANCOVA non-significant	X

V. Part II: Testing on I-CONECT dataset

1. Dataset

Group Session Triad	File Integrity Check	Audio time sync
	Error types • missing file: triad not complete • parse_error: file name error • duplicate: multiple file in one group or $\max(\Delta[p]) < 1\epsilon - 3, \Delta\text{dur}| \leq 0.05s$ • length mismatch: $|\Delta\text{dur}| > 1.0s$ • empty: $\text{dur} < 0.5s$	• cross-correlate loudest 20s window of staff / ppt only file against full conversation padded window • sync_low_corr error raised

11,849 triads (89.9%) are clean without errors.

2. Segmentation, Validation

Participant (ppt only)	rVAD for speech segment detection	
Researcher (staff only)	1. pyannote for speaker diarization	2. mask staff only segments, run rVAD
Full Conversation	30 Week 1 Day 3 VC files annotated manually for automation validation (in progress)	

Staff only audio has participant voice leakage issues, sometimes not detectable by rVAD (for example staff utterance sits at −34.2 dB RMS while the adjacent participant bleed sits at −33.8 dB) so pyannote is used to mask staff only segments. rVAD allows uniform pause detection on diarized audio. pyannote is ran on NYU HPC. Tested on one participant weekly call 138 clips.

Type	Recall	Precision	F1
Sensitive rVAD (v0-3)	94.3 %	86.8 %	0.904
Strict rVAD (v0-9)	77.6 %	95.6 %	0.856
Pyannote	98.5 %	88.8 % (due to pause detection error)	0.934

3. Audio to network transformation

Following Tang et al. (2022), we take the week 1 day 3 of video chats for each participant. Of 63 participants with a Week 1 Day 3 recording, 3 are excluded by the validation pass, leaving: 32 MCI, 28 NC audio.

Then the audio files are transformed to network traffic using RouterSense wireguard VPN. For network traffic transformation jitsi video conference is simulated on ubuntu virtual machine with ppt only audio as the microphone input. Total 1828.4 min of audio is converted to network traffic, and now data analysis is in progress

References

[1] Chou et al. (2024), Mueller (2024), Szatloczki et al. (2015), Chakraborty et al. (2020).
[2] Zhang et al. (2025), Mueller (2024), Szatloczki et al. (2015), Chakraborty et al. (2023), Saeedi et al. (2024)
[3] Parlak et al. (2023)
[4] Dodge HH, Yu K, Wu C-Y, et al. (2024)
[5] Tang F, Chen J, Dodge HH, Zhou J. (2022)
[6] Mahmood R, et al. (2026)

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics and NYU-KAIST Global Innovation and Research Institute.

Engineering Health



Zhaorong Zeng
Chemical and Biomolecular Engineering, B.S.
Spring 2027
NYU Tandon

Design and Prototyping of an Integrated NPLIN Control System for Laser-Induced Crystallization in Supersaturated Solutions

Crystallization is a critical separation and purification process in pharmaceutical and chemical manufacturing, yet controlling when and where nucleation occurs remains a fundamental challenge. Non-photochemical laser-induced nucleation (NPLIN) offers a promising solution by using pulsed laser light to trigger crystallization in supersaturated solutions on demand. Under the leading impurity-heating (IH) mechanism, colloidal nanoparticles absorb laser energy, generate transient vapor nanobubbles, and initiate nucleation at the bubble-solution interface. However, optimizing nanoparticle concentration for reliable crystal yield and developing reproducible experimental platforms remain open challenges.

This project addressed both through two parallel efforts. First, Fe_3O_4 nanoparticles were dispersed in supersaturated KNO_3 (S = 1.06) at three particle number densities — 3×109, 6×109, and 1.2×1010 particles/mL — and tested in a microfluidic capillary platform under fixed laser conditions (50.32 mW/cm², 1064 nm, cold plate 8°C, flow rate 275–300 mL/min). Second, a prototype automated NPLIN control system was developed, consisting of two software applications: one to simultaneously regulate cold plate temperature, flow rate, and laser intensity, and one to automate temperature ramping of supersaturated glycine solutions during preparation.

Crystal yield and size increased with nanoparticle concentration, with an intermediate concentration (~75 µL stock per 40 mL solution) yielding crystals of optimal size and controllability. Higher concentrations produced crystals too large for consistent microfluidic handling. The automated temperature ramping application significantly reduced premature spontaneous nucleation during glycine solution preparation, a persistent challenge in manual protocols.

These results demonstrate that nanoparticle concentration is a tunable parameter for controlling NPLIN crystal output, and that automated system control meaningfully improves experimental reproducibility. The prototype platform developed here represents a step toward a fully integrated, programmable NPLIN system with broad applicability to controlled crystallization in pharmaceutical manufacturing.

NYU TANDON SCHOOL OF ENGINEERING

Prototyping of an Integrated NPLIN Platform: Control System and Effect of Nanoparticle Concentration

Zhaorong Zeng[1], Sammuel Xu[2], Ulises Gutierrez-Nunez[3], Bruce Garetz* and Ryan Hartman*
Department of Chemical and Biomolecular Engineering, NYU Tandon School of Engineering, Brooklyn NY

Background

Crystallization is a fundamental process in pharmaceutical and chemical manufacturing yet controlling nucleation parameters of new crystals reliably remains a significant challenge. Non-photochemical laser-induced nucleation (NPLIN) offers control over crystallization using pulsed laser light, with colloidal nanoparticles playing a central role under the one of the proposed mechanisms underlying NPLIN; impurity-heating mechanism [1] [2]. In this mechanism a nanoparticle absorbs laser energy from a pulse, forms vapor nano bubble and initiates nucleation at the bubble-solution interface.

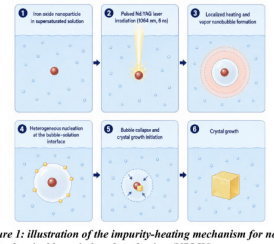


Figure 1: illustration of the impurity-heating mechanism for non-photochemical laser-induced nucleation (NPLIN)

Objective

This project optimized Fe_3O_4 nanoparticle concentration and laser intensity for a microfluidic capillary-based system for NPLIN of a potassium nitrate solution KNO_3 with a super saturation rate of S = 1.06 (holding 6% more of solute than water can hold at a temperature of 20° Celsius); moreover, we developed prototype software to automate experimental control: flow rate, temperature of the solution & laser intensity which are all parameters that affect NPLIN.

Methodology

Aqueous supersaturated solutions of Potassium Nitrate were prepared with a range of Fe_3O_4 dispersions (51, 75, 101, 202 µL of 0.17mg/mL per 40 mL) added to the solutions. We tested with a new protocol of water bath ramping for supersaturated solutions preparation to prevent spontaneous crystallization. Using a temperature-controlled chiller machine and a custom-made software. Different incident laser intensities were tested in range from 45 to 60 mW/cm². We used a cold Peltier module to generate a thermo-cooling effect to increase super saturation state. Flow rate was delivered by a Piston pump and a set of tubing connected to the capillary microfluidic device. A custom-made software was built using Python and serial communication programs to connect with the lab machines to regulate thermo-cooling, flow rate, and laser intensity. All controlled by user end on computer.

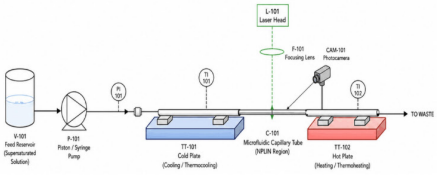


Figure 2: Schematic diagram of the microfluidic NPLIN experimental system.

Results

We stablished an optimal temperature of 3°C and found an ideal flow rate of 0.300 mL/min. The 75 µL of the dispersed nanoparticles addition in the 40 mL of the KNO3 solution and 55 mW/cm² of incident laser intensity were found to be the optimal, maximizing crystal yield while preventing the capillary system from clogging caused by oversized crystals at higher nanoparticle concentrations and laser intensities. The system improved reproducibility and reduced spontaneous nucleation during solution preparation.

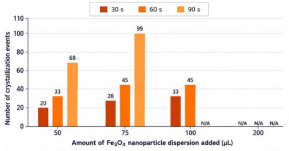


Figure 3: Effect of Fe_3O_4 nanoparticle amount on crystallization performance.

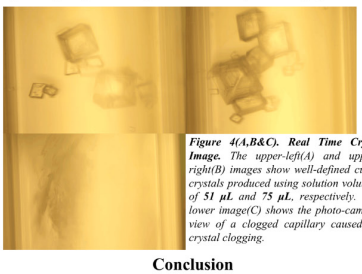


Figure 4(A,B&C). Real Time Crystal Image. *The upper-left(A) and upper-right(B) images show well-defined cubic crystals produced using solution volumes of **51 µL** and **75 µL**, respectively. The lower image(C) shows the photo-camera view of a clogged capillary caused by crystal clogging.*

Conclusion

The prototype control system integrates regulation of temperature, flow rate, and laser intensity into a single platform which significantly improves experimental reproducibility enabling more consistent NPLIN conditions and reliable data recording. These results establish that 75 µL Fe_3O_4 and 55 mW/cm² are the optimal conditions for microfluidic NPLIN of KNO_3, balancing crystal yield against capillary clogging.

Works Cited

[1] Garetz, B. A.; Hartman, R. L. 25+ Years of Research on Nonphotochemical Laser-Induced Nucleation (NPLIN). *Cryst. Growth Des.* ASAP. DOI: 10.1021/acs.cgd.5c00266.
[2] Ndukwe-Ajala, K. F.; Subrin, J. M.; Garetz, B. A.; Hartman, R. L. Microfluidic Laser-Induced Nucleation of Iron(II,III) Oxide Nanoparticle-Doped Supersaturated Aqueous KCl Solutions. Cryst. Growth Des. 2024, 24 (20), 8370–8380. DOI: 10.1021/acs.cgd.4c00805
[3] Cui, P.; Kotede, V.; van Tooren, R. P. Y.; Nagalingam, N.; Wang, R.; Yin, Q.; van der Heijden, A. E. D. M.; Kramer, H. J. M.; Eral, H. B. Influence of Nanoparticle Properties on Non-Photochemical Laser-Induced Nucleation. *Cryst. Growth Des.* **2026**, 26 (1), 144–154. DOI: 10.1021/acs.cgd.5c01080.

Contact: Corresponding authors: *bgaretz@nyu.edu *ryan.hartman@nyu.edu 1: zz10523@nyu.edu 2: mx2177@nyu.edu 3: ug2058@nyu.edu



Systems Engineering & Complex Decision-Making

The research and courses in systems engineering and data-driven decision-making tackle real-world challenges, creating secure and effective solutions for everything from global financial systems and supply chains to digital healthcare, urban transportation, and critical infrastructure.

By combining advanced urban data modeling methods and techniques with systems engineering principles and exposure to a variety of practical applications, we're building a new generation of engineers and interconnected systems that are resilient, efficient, and ready for the future.

Our areas of focus include: Intelligent digital health infrastructure, Energy systems Complex, dynamical systems, Decision analytics, Financial risk and Industrial engineering

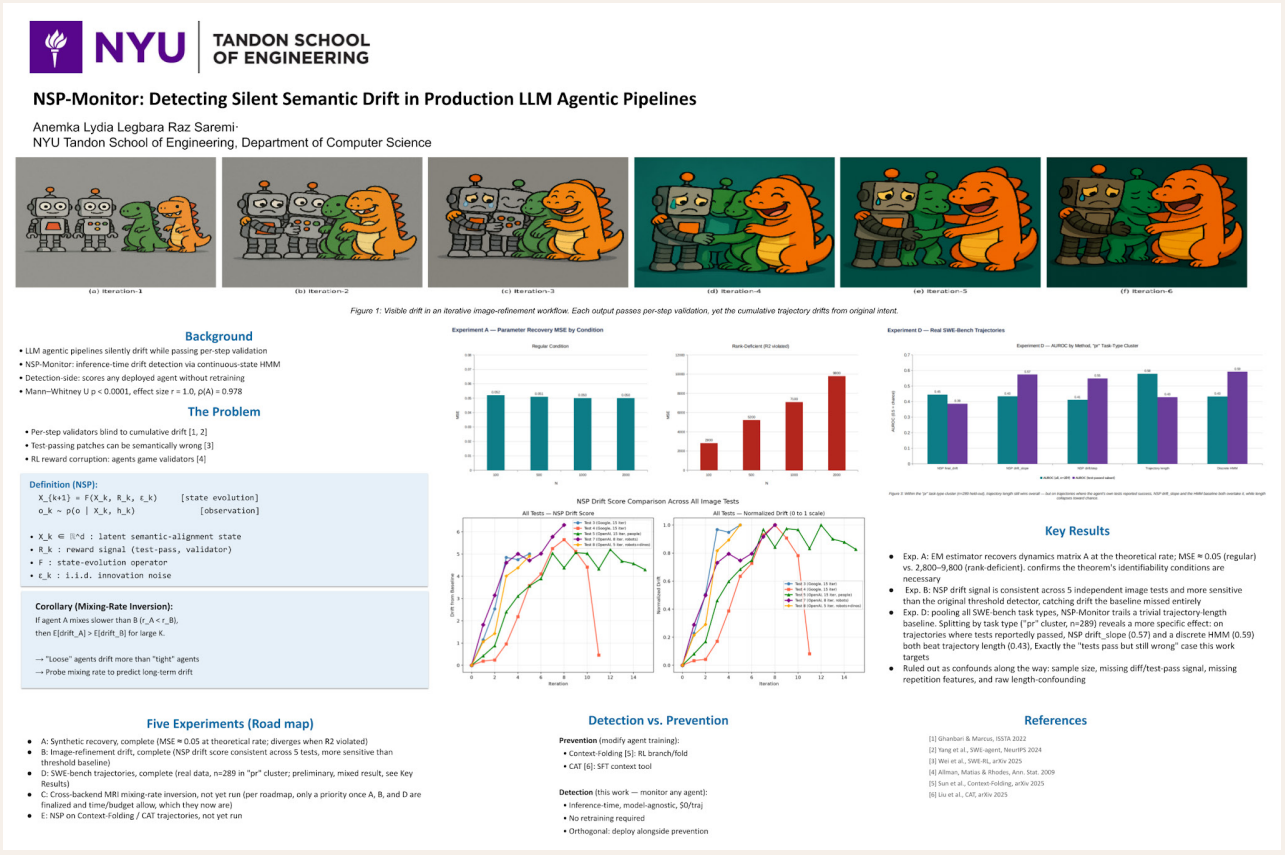
Systems Engineering & Complex Decision-Making



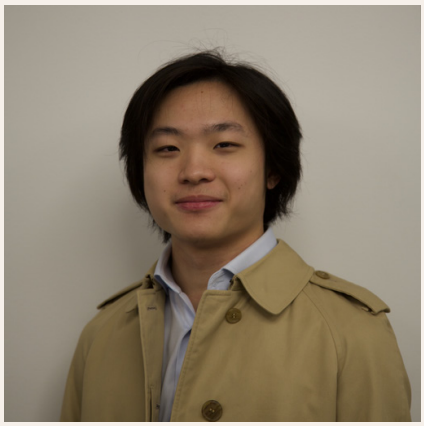
Anemka Lydia Legbara
Computer Science, B.S.
Spring 2027
NYU Tandon

NSP-Monitor: Detecting Silent Semantic Drift in Production LLM Agentic Pipelines

Iterative LLM agentic pipelines, where an agent operates sequentially on its own previous output, exhibit a critical class of silent failures: individual steps pass validation, yet the cumulative trajectory drifts from the user's original intent. This work introduces NSP-Monitor, an inference-time drift detection tool grounded in the Neural-trace Semantic-state Process (NSP), a continuous-state hidden Markov model that formalizes how an agent's latent semantic-alignment state evolves under potentially corrupted reward signals. Unlike recent prevention-side approaches such as Context-Folding and CAT that modify agent training to manage context, NSP-Monitor operates on the detection side: it scores trajectories produced by any deployed agent, without retraining, and flags those where cumulative drift exceeds a calibrated threshold. The framework is validated through five experiments: synthetic parameter recovery confirming estimator convergence at the Cramér–Rao rate; drift detection on SWE-bench Verified trajectories, our primary industrial case study, where we compare NSP-Monitor against simple and model-based baselines; human-correlated drift measurement on iterative image-refinement workflows; cross-backend mixing-rate inversion demonstrating that per-step edit style predicts long-term drift behavior; and a complementarity study applying NSP-Monitor to Context-Folding and CAT trajectories to test whether prevention-side approaches eliminate the need for detection. Preliminary results on 40 multi-turn trajectories (20 stable, 20 exhibiting induced drift) show that the NSP drift-slope score separates the two groups with a Mann–Whitney U $p < 0.0001$, a rank-biserial effect size of $r = 1.0$, and a fitted spectral radius of $\rho(A) = 0.978$ confirming stable but near-critical mixing dynamics. The EM estimator converges within 30 iterations and recovers dynamics-matrix parameters with MSE < 0.05 under the identifiability conditions specified by the regularity assumptions, while deliberately violating the full-rank Fisher information condition (R2) causes estimation to diverge, empirically confirming the necessity of the theorem's assumptions.



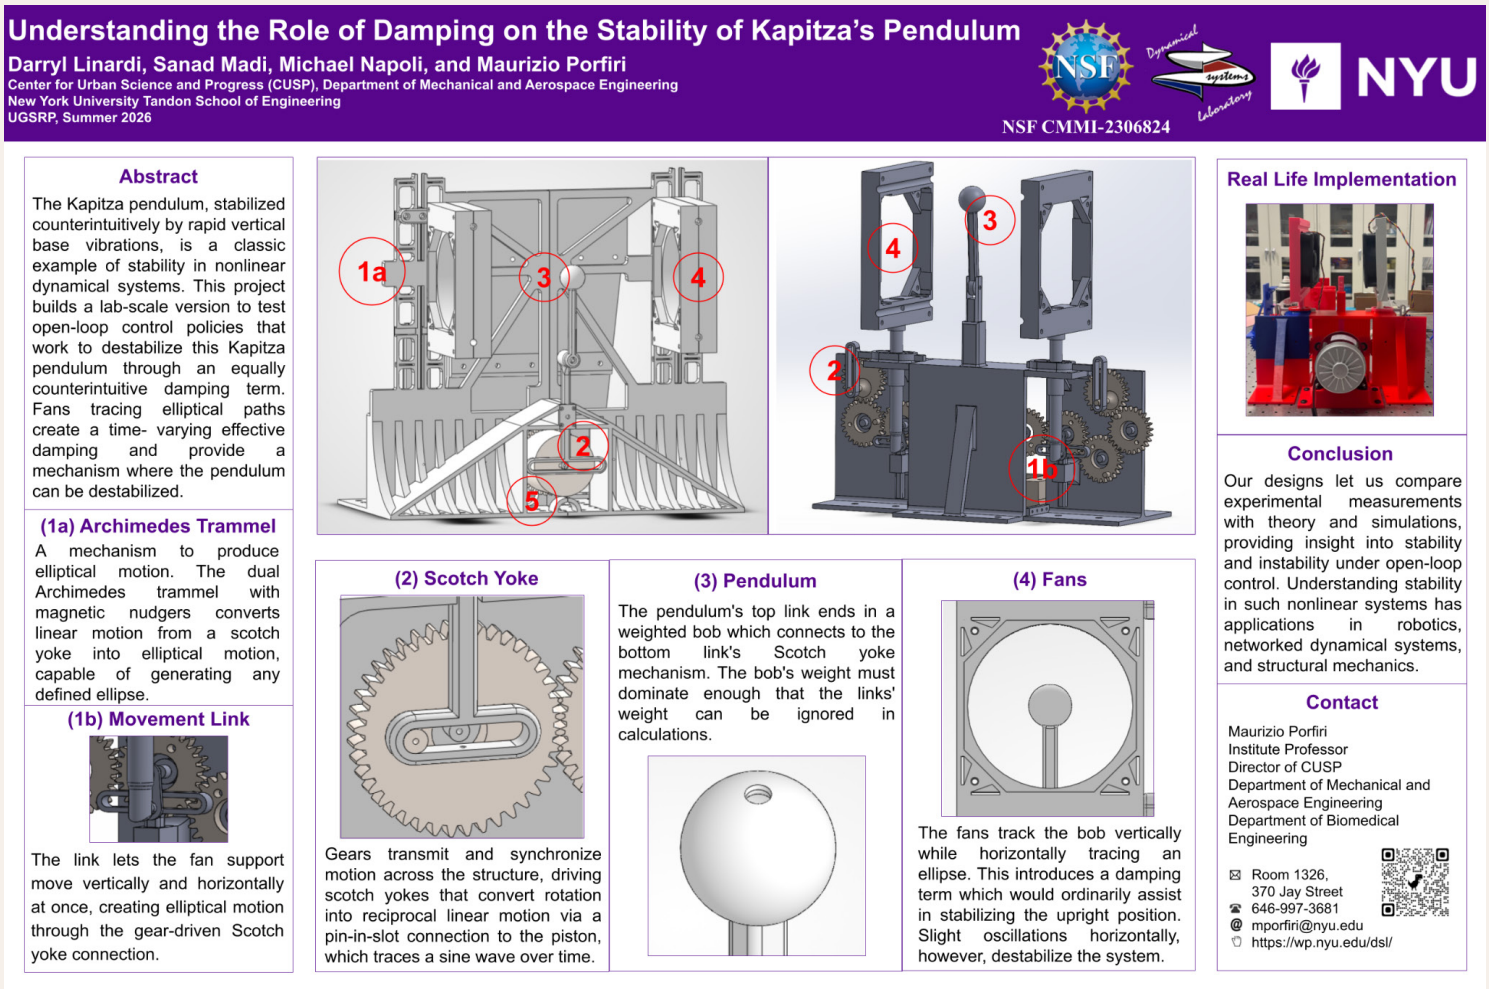
Systems Engineering & Complex Decision-Making



Darryl Linardi
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon

Understanding the Role of Damping on the Stability of Kapitza's Pendulum

Stability, control, and chaos are among the most fascinating and important ideas in the study of dynamical systems. The Kapitza pendulum is a well-documented example of a nonlinear dynamical system that somewhat counterintuitively stabilizes itself via rapid vertical vibration at its base, making the Kapitza pendulum a particularly useful dynamical system to study. The setup we built consists of a pendulum flanked on either side by fans that track vertically but also have a periodical horizontal motion, resulting in fans that trace ellipses. The frequency and amplitude can be systematically varied to adjust the stability region of the pendulum. The fans introduce a damping coefficient which makes the upright position more stable but if the fans wiggle even slightly, the system immediately destabilizes. Experimentally measured data of the pendulum's motion will be compared with theoretical predictions and numerical simulations to identify regimes of stability and instability under open-loop control. Understanding the emergence of stability in complex nonlinear systems has applications in robotics, networked dynamical systems, and structural mechanics.



Systems Engineering & Complex Decision-Making



Evan Dworkin
Electrical Engineering and Computer Engineering (dual major), B.S.
Fall 2026
NYU Tandon

Open-Source Chiplet Electrical Interconnect Modeling

The standard method of computer chip manufacturing over the last few decades has been to package the chip as one monolithic design on a single die. However, chiplet technologies have emerged in recent years as a new frontier for improving computing performance. A chiplet-based approach divides the functions of the chip into discrete units (chiplets) that can be manufactured each with higher yield on the optimal process for that chiplet. In order to communicate with one another these chiplets must have electrical interconnects, which has sparked the need for chiplet-specific design tools to aid computer architects. This project, Chiplet Link IP Generation (CLIPGen) fills this gap with an easy-to-configure electrical link design search space exploration. The chiplet link characterization procedure can be executed through proprietary, closed source tools (Cadence Liberate) or through a completely open-source pipeline (CharLib, ngspice, and FreePDK45). CLIPGen brings fast, reliable, and accurate chiplet link modeling to aid computer architects in making critical early-design choices. The open-source pipeline puts this power in the hands of the general public by lowering the barrier of entry to advanced integrated circuit design.

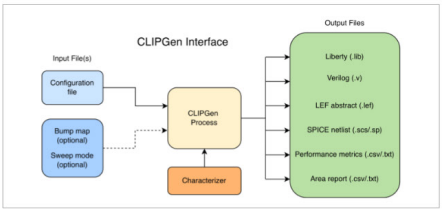
NYU TANDON SCHOOL OF ENGINEERING

Open-Sourcing a Chiplet Electrical Interconnect Model

Evan Dworkin Under the guidance of: Zhengping Zhu, Professor Austin Rovinski
New York University Tandon School of Engineering, Electrical and Computer Engineering Department

Abstract

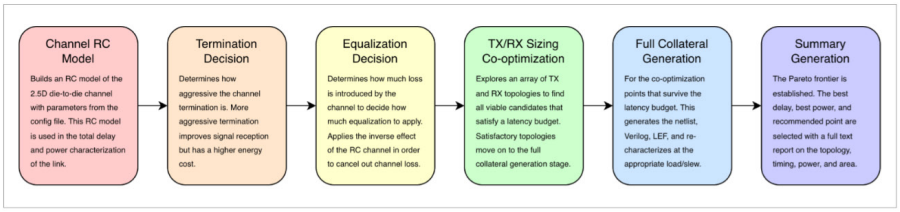
The standard method of computer chip manufacturing over the last few decades has been to package the chip as one monolithic design on a single die. However, in recent years, chiplet technologies have emerged as a new frontier for improving computing performance. This project, Chiplet Link IP Generation (CLIPGen)[1] fills this gap with an easy-to-configure electrical link design search space exploration. The chiplet link characterization procedure can be executed through proprietary, closed source tools (Cadence Liberate, TSMC65 PDK) or through a completely open-source pipeline (CharLib, ngspice, and FreePDK45). CLIPGen performs fast, reliable, and accurate chiplet link modeling to aid computer architects in making critical early-design choices for chiplet communication. The open-source pipeline puts this power in the hands of the general public by lowering the barrier of entry to advanced integrated circuit design.



Project Background

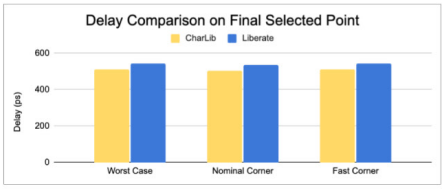
A chiplet-based approach to computer chip design divides the functions of the chip into discrete units (chiplets) that can be manufactured each with higher yield on the optimal process for that chiplet. The use of chiplets can decrease production times and lower manufacturing costs. However, in order to communicate with one another these chiplets must have electrical interconnects, which has sparked the need for chiplet-specific design tools to aid computer architects. Currently existing tools are either too high-level to be useful in making design decisions or model the interconnect too deeply, requiring extensive expertise from the architect. CLIPGen is designed to be user-friendly to the less experienced architect while still offering valuable and customizable data to the expert. The entire process can be controlled through a single configuration file with knobs for:

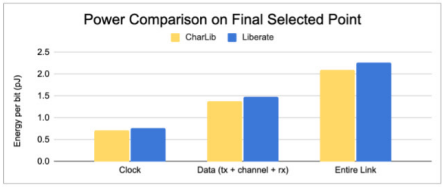
- interconnect reach (1-50 mm)
- bump pitch (1-130 µm)
- data rate (1-48 Gb/s)
- lane count (x8, x16, x32, x64)
- package type (silicon / organic)
- PDK (any PDK can be dropped in)
- characterizer backend (Cadence Liberate/Spectre or CharLib/ngspice)



Open-Source Objective

As processing technologies get more advanced, the barrier to entry of chip design gets higher and higher. The expertise required to use modern design tools combined with the extreme cost of software licenses limits aspiring and amateur chip designers. In recent years there have been several pushes to make chip design tools more accessible with open-source software (i.e. OpenROAD, CharLib), PDKs (i.e. FreePDK, SkyWaterPDK), and simulators (i.e. ngspice, Xyce). CLIPGen is designed to work with both the industry-standard commercial tools as well as with open-source tools to be accessible to all.





Results Comparison

Both the commercial (Liberate) and open-source (CharLib) characterizer backends agree on the final recommend design point parameters. CharLib characterizes the delay at a consistent 7% faster than Liberate and characterizes the power at a consistent 8% less than Liberate.

Conclusions

The open-source backend (CharLib + ngspice) has been able to match the output of the commercial backend (Cadence Liberate + Spectre) to within 8%. In terms of design decisions, both backends produced the same TX topology (6 stages, 2.83 stage ratio, 1.5 beta ratio) and the same RX topology (0.4µm / 0.6µm preamp sizing, 0.6µm / 0.9µm buffer sizing). The discrepancy between the delay and power characterization boils down to differences in the simulator engines. Since Spectre is a closed-source, proprietary simulator it is impossible to know what it does under the hood to generate its numbers. Therefore, it is impossible to know what issue underpins the discrepancy. To determine which simulator is correct would require real fabrication and testing, which is impractical for now. For the time being, the numbers agree well enough to serve CLIPGen's purpose of aiding architects in making early design decisions about their chip, chiplets, and the interconnects between them.

Future Work

The CLIPGen project addresses both the need for high quality chiplet modeling tools and the accessibility crisis of advanced semiconductor manufacturing. Further work can be done in improving the model's accuracy and verifying results with full 3D electromagnetic simulations. Fabrication and rigorous testing of CLIPGen-generated links can be the final step in confirming the model's accuracy. This verification can be used for both commercial tools and open-source tools as we strive for full parity between the two. As it stands, CLIPGen is a valuable tool for computer architects making early design stage decisions with 2.5D packaging at up to 48Gb/s. Future additions to the framework include 3D packaging support and support for inductance-aware interconnects that support faster data rates.

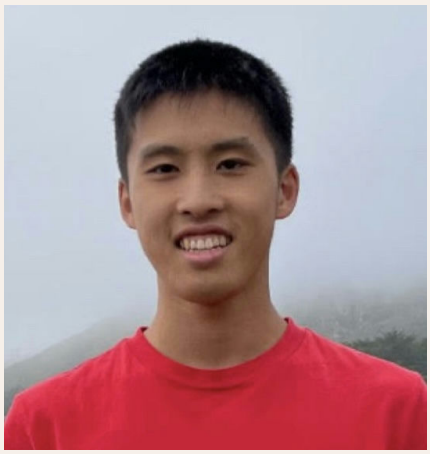
References

[1] Zhu, Zheng and Austin Rovinski. "CLIPGen: A Chiplet Link IP Modeling and Generation Framework for 2.5D Architecture Exploration." ArXiv abs/2605.27757 (2026): n. pag.

Acknowledgement

The author thanks the NYU Tandon School of Engineering's Office of Undergraduate Academics and the NYU Tandon GLASS Honors Program for generous funding of the project. Additional thanks to Professor Austin Rovinski and Zhengping Zhu for your guidance and insight.

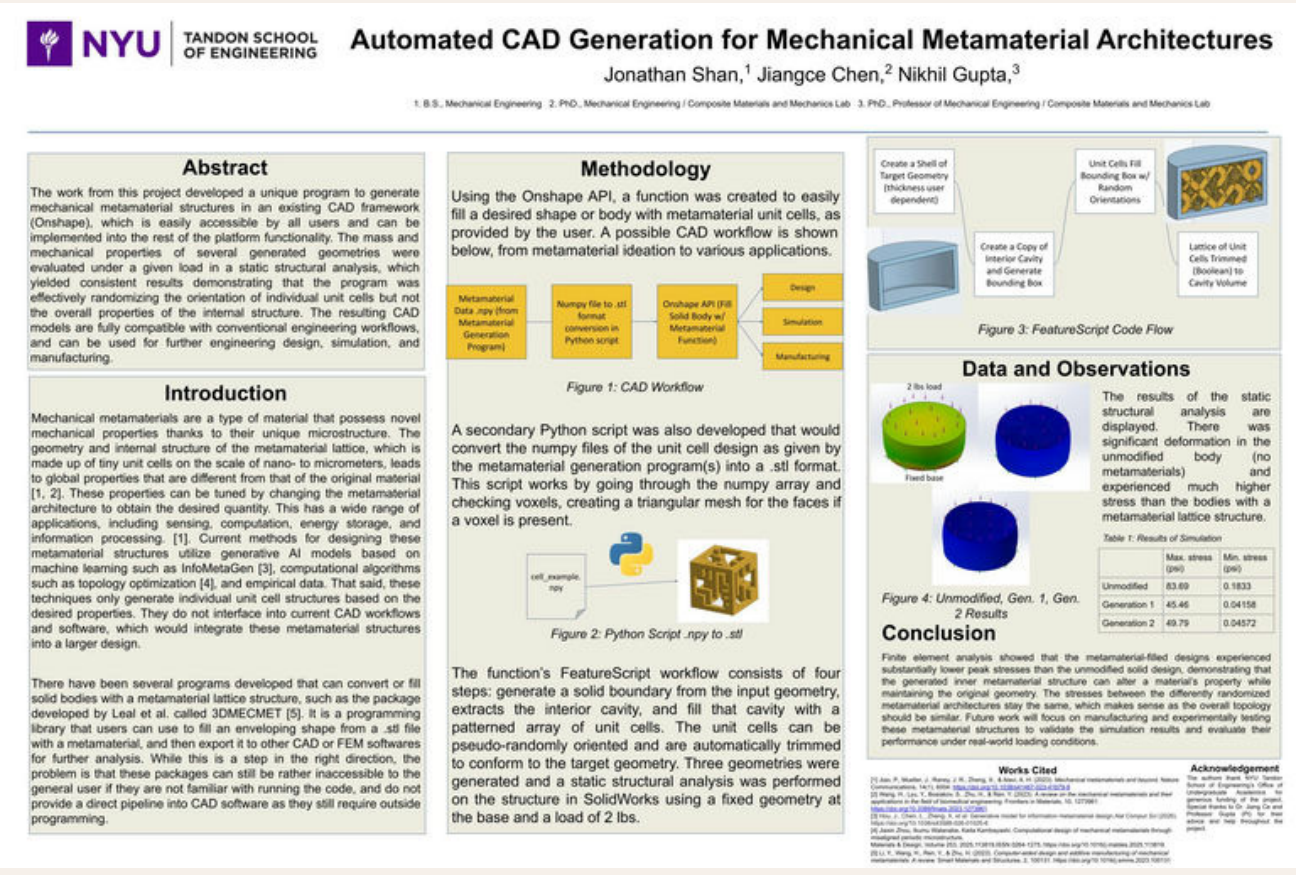
Systems Engineering & Complex Decision-Making



Jonathan Shan
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon

An Automated CAD Generation Pipeline for Mechanical Metamaterial Design

Mechanical metamaterials are materials that possess unique mechanical properties thanks to their geometry and internal structure. These properties can be tuned by changing the material architecture to obtain the desired quantity. This has a wide range of applications, including sensing, computation, energy storage, and information processing. Current methods for designing these metamaterial structures utilize computational modeling and special algorithms that output a unique unit cell lattice structure, but they lack integration into current CAD interfaces and workflows. There have been several attempts by researchers who have developed programs/packages that can fill an enveloping shape from a CAD .stl file with a certain metamaterial, but they are still rather inaccessible to the general user and do not provide a full pipeline into CAD software as they still require outside programming. The work from this project developed a functional metamaterial generation program on an existing CAD framework (Onshape) which is easily accessible by all users and can be implemented into the rest of the platform functionality. A unique function was written using Onshape's FeatureScript which can be called to directly generate a metamaterial structure in an Onshape part studio, and then be manipulated with other CAD features. The function populates arbitrary three-dimensional volumes with user-defined unit cells, applies randomized orientations to improve structural diversity, trims cells to the boundary of the object/volume, and generates a continuous solid boundary for fabrication. The mass and mechanical properties of several generated geometries were evaluated under a given load in a static structural analysis, which yielded consistent results demonstrating that the program was effectively randomizing the orientation of individual unit cells but not the overall properties of the internal structure. The resulting CAD models are fully compatible with conventional engineering workflows, and can be used for further engineering design, simulation, and manufacturing. This eliminates the need for extensive manual modeling while enabling rapid design and evaluation of complex metamaterial structures.



Systems Engineering & Complex Decision-Making

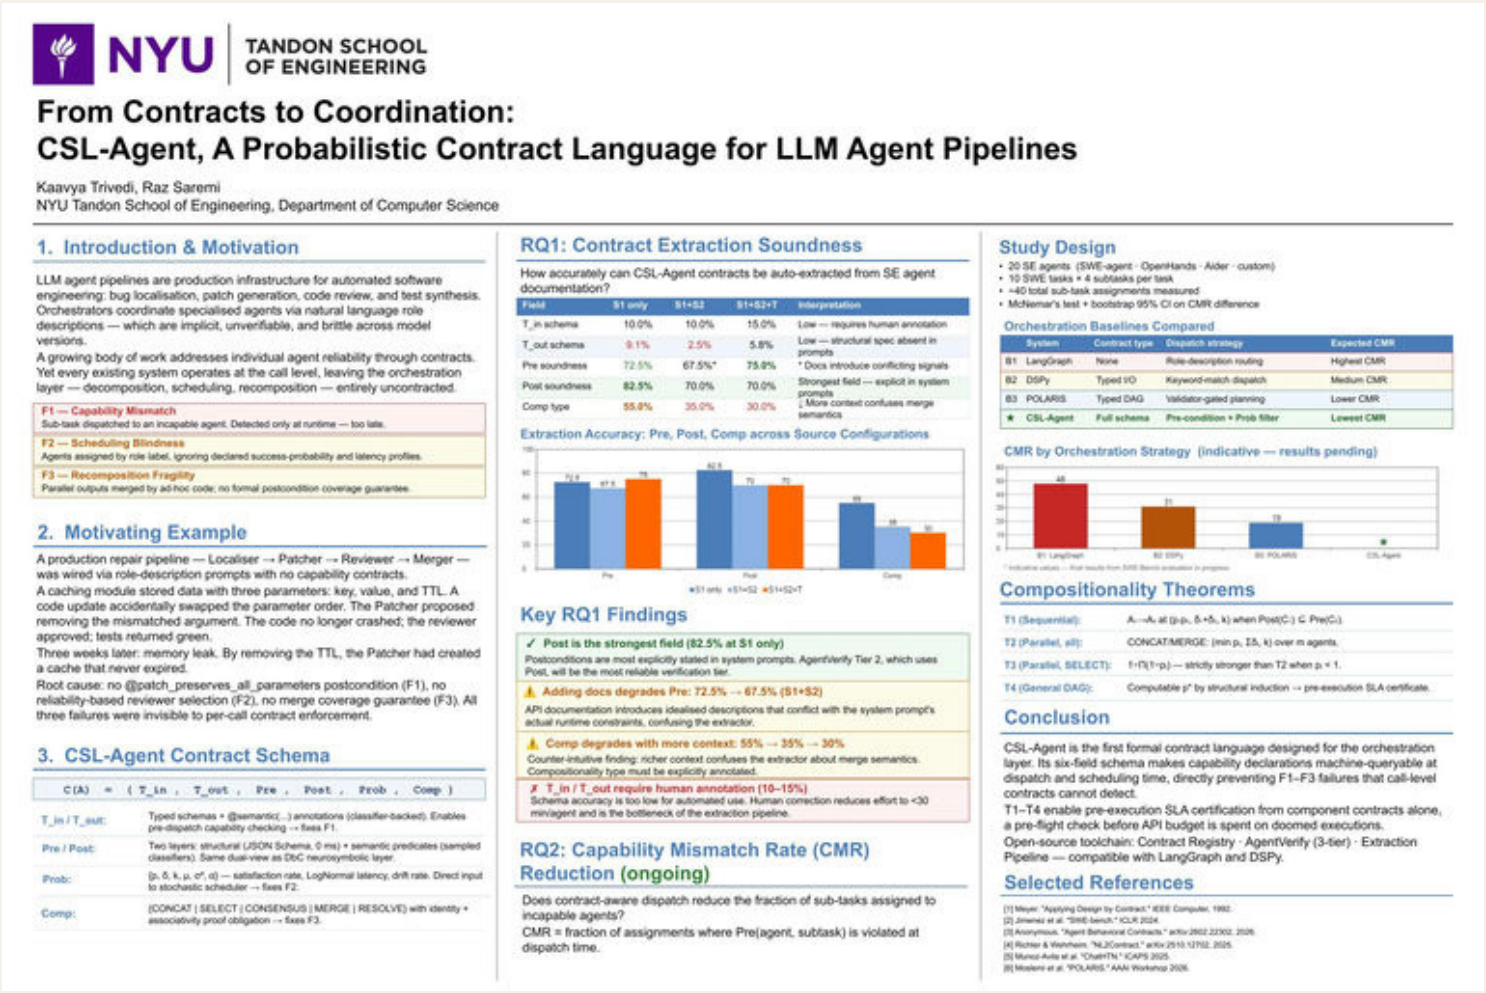


Kaavya Trivedi
Computer Science, B.S.
Spring 2029
NYU Tandon
Thompson-Bartlett Fellow

CSL-Agent: Contract-Based Specification and Runtime Verification for Multi-Agent Software Engineering Pipelines

The deployment of LLM-based agents in multi-agent software engineering (SE) pipelines has introduced a critical reliability gap: agents are routinely dispatched to tasks their capabilities cannot support, producing silent failures with no formal mechanism for detection or prevention. This paper presents CSL-Agent, a contract specification language and runtime verification framework that derives typed behavioral contracts, encoding input/output schemas, preconditions, and postconditions, directly from SE agent documentation and system prompts. Unlike prompt-engineering and fine-tuning approaches that embed behavioral expectations implicitly, CSL-Agent makes those expectations explicit and machine-checkable: contracts are extracted once and enforced at every dispatch decision without modifying the underlying agent.

We address the following research question grounded in a corpus of 20 real-world SE agents drawn from SWE-agent, OpenHands, Aider, and more: How accurately can contracts be extracted from agent documentation across four information configurations, system prompt alone, system prompt with documentation, with execution traces, and with human annotation?



Systems Engineering & Complex Decision-Making



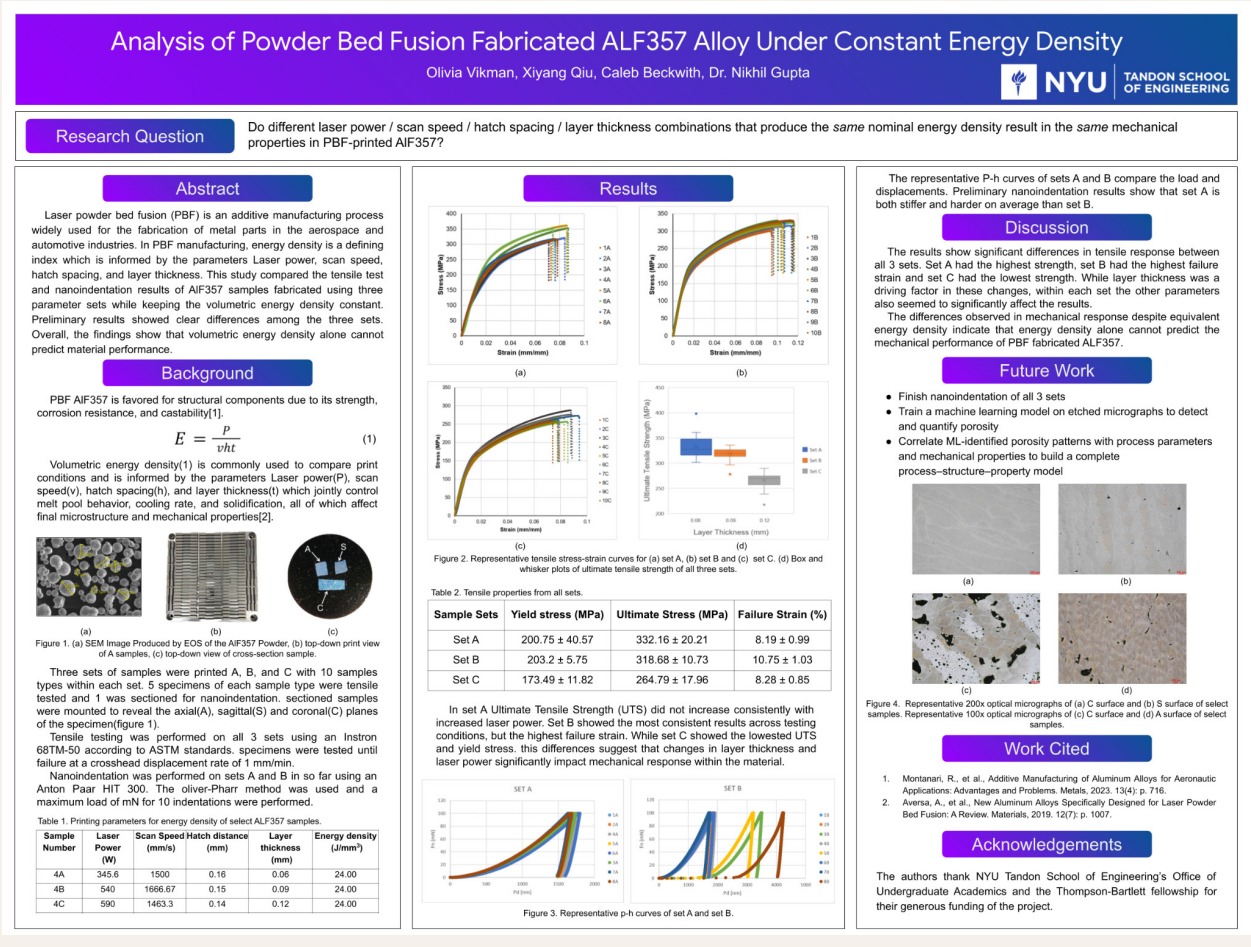
Olivia Vikman
Mechanical Engineering, B.S.
Spring 2028
NYU Tandon
Thompson-Bartlett Fellow



Xiyang Qiu
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon

Analysis of Laser Powder Bed Fusion Fabricated Aluminum Alloy AlF357 under Constant Energy Density

Laser powder bed fusion (PBF) is an additive manufacturing process widely used for the fabrication of metal parts in the aerospace and automotive industries. In PBF manufacturing, energy density is a defining index by which the properties of manufactured parts can be characterized. However, specimens with the same energy density can still show different mechanical properties. This is because energy density is informed by the parameters Laser power, scan speed, hatch spacing, and layer thickness. This study investigates 3 sets of PBF fabricated AlF357 aluminum alloy at a constant volumetric energy density of 24 J/mm^3 while varying the laser power, scan speed, hatching speed and layer thickness. Cross-sectional specimens were prepared for optical microscopy and nanoindentation. Tensile testing was used to determine mechanical properties according to ASTM standards. Optical microscopy revealed orientation dependent microstructural features and printing defects. The tensile results showed clear differences among the three sets. Set A had the thinnest layers and produced the highest ultimate tensile strength, while Set C had the thickest layers and showed the lowest overall strength. Overall, the findings show that volumetric energy density alone cannot predict material performance and individual processing parameters may have measurable effect on the properties of the printed part.



Analysis of Powder Bed Fusion Fabricated ALF357 Alloy Under Constant Energy Density

Olivia Vikman, Xiyang Qiu, Caleb Beckwith, Dr. Nikhil Gupta

NYU TANDON SCHOOL OF ENGINEERING

Research Question

Do different laser power / scan speed / hatch spacing / layer thickness combinations that produce the *same* nominal energy density result in the *same* mechanical properties in PBF-printed AlF357?

Abstract

Laser powder bed fusion (PBF) is an additive manufacturing process widely used for the fabrication of metal parts in the aerospace and automotive industries. In PBF manufacturing, energy density is a defining index which is informed by the parameters Laser power, scan speed, hatch spacing, and layer thickness. This study compared the tensile test and nanoindentation results of AlF357 samples fabricated using three parameter sets while keeping the volumetric energy density constant. Preliminary results showed clear differences among the three sets. Overall, the findings show that volumetric energy density alone cannot predict material performance.

Background

PBF AlF357 is favored for structural components due to its strength, corrosion resistance, and castability[1].

$$E = \frac{P}{vht} \quad (1)$$

Volumetric energy density(1) is commonly used to compare print conditions and is informed by the parameters Laser power(P), scan speed(v), hatch spacing(h), and layer thickness(t) which jointly control melt pool behavior, cooling rate, and solidification, all of which affect final microstructure and mechanical properties[2].

(a) (b) (c)

Figure 1. (a) SEM Image Produced by EOS of the AlF357 Powder, (b) top-down print view of A samples, (c) top-down view of cross-section sample.

Three sets of samples were printed A, B, and C with 10 samples types within each set. 5 specimens of each sample type were tensile tested and 1 was sectioned for nanoindentation. sectioned samples were mounted to reveal the axial(A), sagittal(S) and coronal(C) planes of the specimen(figure 1).

Tensile testing was performed on all 3 sets using an Instron 68TM-50 according to ASTM standards. specimens were tested until failure at a crosshead displacement rate of 1 mm/min.

Nanoindentation was performed on sets A and B in so far using an Anton Paar HIT 300. The oliver-Pharr method was used and a maximum load of mN for 10 indentations were performed.

Table 1. Printing parameters for energy density of select ALF357 samples.

Sample Number	Laser Power (W)	Scan Speed (mm/s)	Hatch distance (mm)	Layer thickness (mm)	Energy density (J/mm^3)
4A	345.6	1500	0.16	0.06	24.00
4B	540	1666.67	0.15	0.09	24.00
4C	590	1463.3	0.14	0.12	24.00

Results

Figure 2. Representative tensile stress-strain curves for (a) set A, (b) set B and (c) set C. (d) Box and whisker plots of ultimate tensile strength of all three sets.

Table 2. Tensile properties from all sets.

Sample Sets	Yield stress (MPa)	Ultimate Stress (MPa)	Failure Strain (%)
Set A	200.75 ± 40.57	332.16 ± 20.21	8.19 ± 0.99
Set B	203.2 ± 5.75	318.68 ± 10.73	10.75 ± 1.03
Set C	173.49 ± 11.82	264.79 ± 17.96	8.28 ± 0.85

In set A Ultimate Tensile Strength (UTS) did not increase consistently with increased laser power. Set B showed the most consistent results across testing conditions, but the highest failure strain. While set C showed the lowested UTS and yield stress. this differences suggest that changes in layer thickness and laser power significantly impact mechanical response within the material.

Figure 3. Representative p-h curves of set A and set B.

The representative P-h curves of sets A and B compare the load and displacements. Preliminary nanoindentation results show that set A is both stiffer and harder on average than set B.

Discussion

The results show significant differences in tensile response between all 3 sets. Set A had the highest strength, set B had the highest failure strain and set C had the lowest strength. While layer thickness was a driving factor in these changes, within each set the other parameters also seemed to significantly affect the results.

The differences observed in mechanical response despite equivalent energy density indicate that energy density alone cannot predict the mechanical performance of PBF fabricated ALF357.

Future Work

- Finish nanoindentation of all 3 sets
- Train a machine learning model on etched micrographs to detect and quantify porosity
- Correlate ML-identified porosity patterns with process parameters and mechanical properties to build a complete process–structure–property model

(a) (b) (c) (d)

Figure 4. Representative 200x optical micrographs of (a) C surface and (b) S surface of select samples. Representative 100x optical micrographs of (c) C surface and (d) A surface of select samples.

Work Cited

1. Montanari, R., et al., Additive Manufacturing of Aluminum Alloys for Aeronautic Applications: Advantages and Problems. Metals, 2023. 13(4): p. 716.
2. Aversa, A., et al., New Aluminum Alloys Specifically Designed for Laser Powder Bed Fusion: A Review. Materials, 2019. 12(7): p. 1007.

Acknowledgements

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics and the Thompson-Bartlett fellowship for their generous funding of the project.

Systems Engineering & Complex Decision-Making



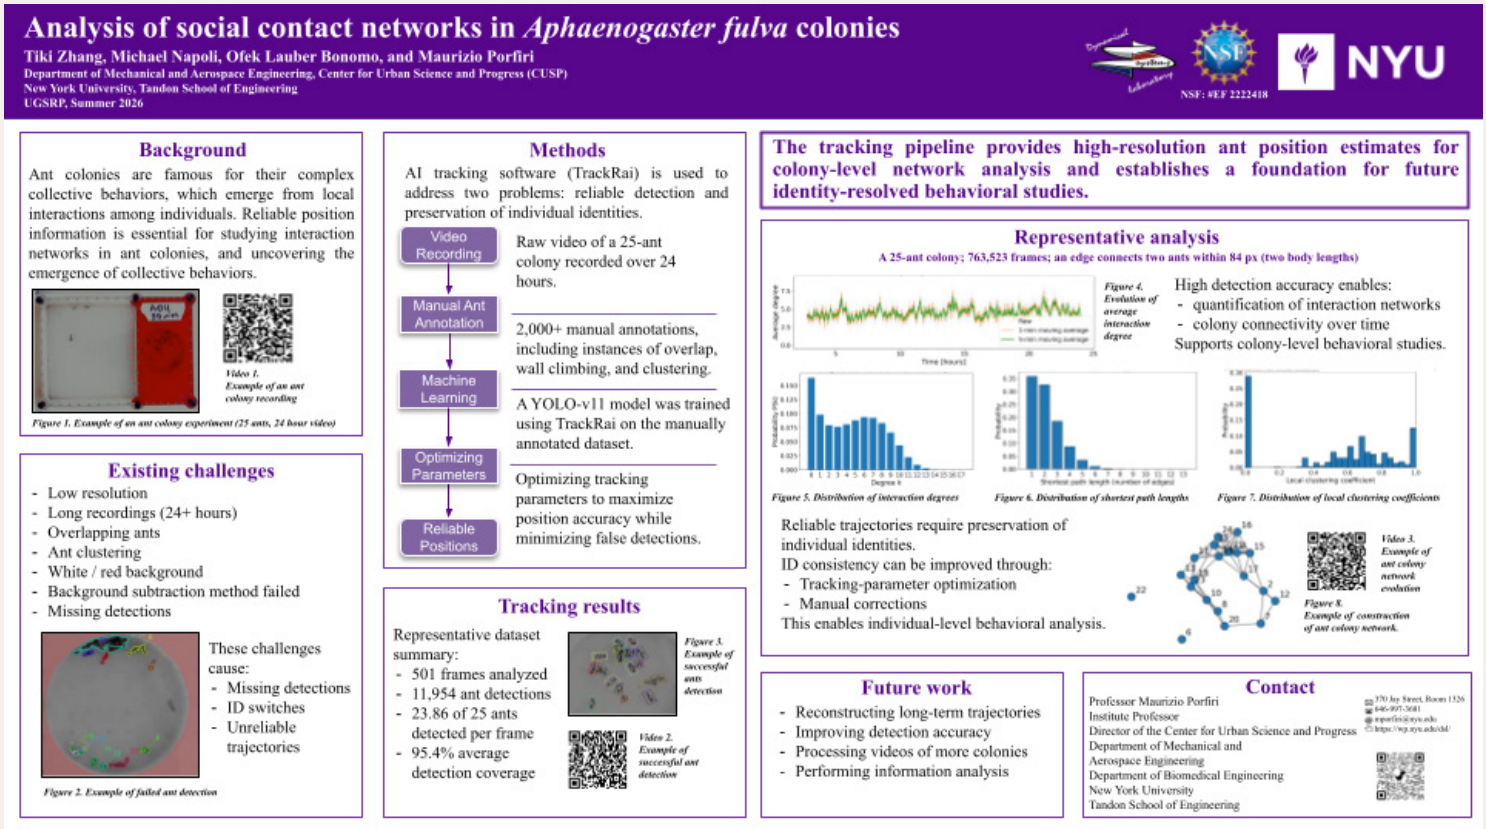
Tiankai Zhang
Mathematics, B.S.
Spring 2028
NYU Tandon

Computational Analysis of Collective Behavior in Ant Colonies

Ant colonies are famous for exhibiting a wide range of collective behaviors - macroscopic behaviors that emerge in the absence of any centralized control - as a function of simple, local interactions between individuals. Due to their complicated temporal and spatial scale, understanding these phenomena requires reliable methods for extracting movement and social information from long-term video recordings. This project utilizes a video-based, computational framework to quantify ant activity and social interactions from an Aphaenogaster fulva ant colony.

Tracking data were processed to reconstruct individual trajectories and identify interactions between ants over time. Activity levels were quantified from movement data, and an integrated interaction network was constructed by accumulating pairwise contacts throughout the observation period. The robustness of activity measurements was evaluated using multiple threshold selection methods.

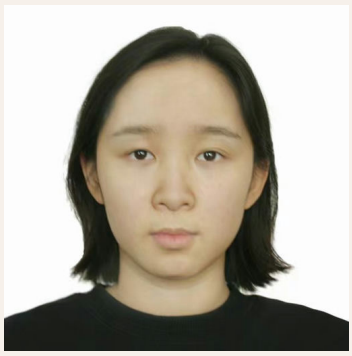
The resulting framework converts raw tracking data into quantitative descriptions of colony behavior, including activity dynamics and long-term interaction networks. Our analyses provide a foundation for studying communication and information transfer underlying collective behavior in social insect colonies, and serves as a first step to quantifying how complex behavioral cues travel across the available network.



Systems Engineering & Complex Decision-Making



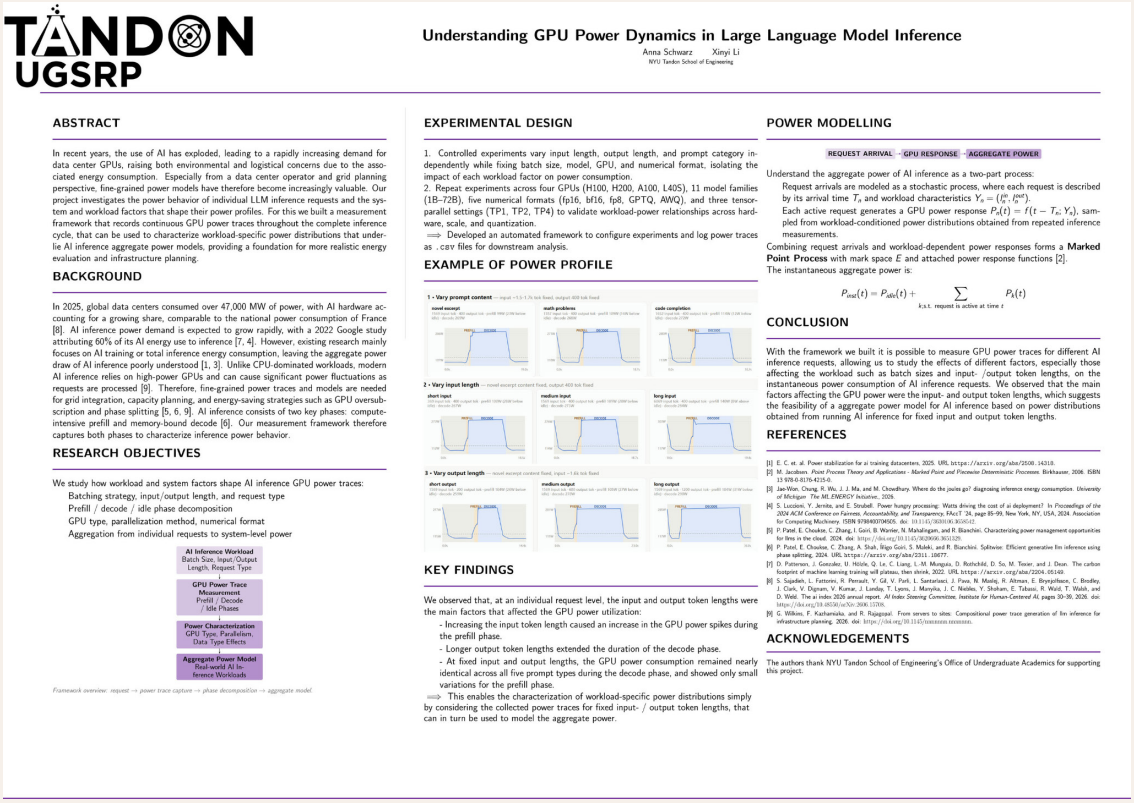
Anna Schwarz
Mathematics, B.S.
Spring 2026
Humboldt University



Xinyi Li
Mathematics, B.S.
Spring 2027
NYU Tandon

Understanding GPU Power Dynamics in Large Language Model Inference

In recent years, the use of AI has exploded, leading to a rapidly increasing demand for data center GPUs, raising both environmental and logistical concerns due to the associated energy consumption. Especially from a data center operator and grid planning perspective, fine-grained power models have therefore become increasingly valuable, for example for grid incorporation, provisioning and capacity planning, all factors that need to be taken into account when trying to make data centers more energy efficient. Although model training is highly energy intensive, the large volume of daily inference requests means that AI inference may account for a substantial share of future AI computing demand. This study therefore investigates the power behavior of individual LLM inference requests and the system and workload factors that shape their power profiles. We developed a measurement framework that records continuous GPU power traces throughout the complete inference cycle, including idle, prefill, and decode phases. Experiments were conducted across a broad range of hardware platforms, language models, prompt categories, batch settings, parallelism strategies, and precision and quantization methods. Our results show that GPU power consumption varies considerably across models, hardware platforms, and serving configurations, while the effect of prompt content is generally much smaller. We also observed that initialization behavior, model optimization, and runtime settings can influence both power characteristics and overall inference efficiency. The collected measurements enable us to characterize workload-specific power distributions rather than relying only on aggregate averages. These distributions further support the development of stochastic models for estimating aggregate power demand in data centers running large-scale AI inference workloads, providing a foundation for more realistic energy evaluation and infrastructure planning.



TANDON UGSRP

Understanding GPU Power Dynamics in Large Language Model Inference

Anna Schwarz Xinyi Li

NYU Tandon School of Engineering

ABSTRACT

In recent years, the use of AI has exploded, leading to a rapidly increasing demand for data center GPUs, raising both environmental and logistical concerns due to the associated energy consumption. Especially from a data center operator and grid planning perspective, fine-grained power models have therefore become increasingly valuable. Our project investigates the power behavior of individual LLM inference requests and the system and workload factors that shape their power profiles. For this we built a measurement framework that records continuous GPU power traces throughout the complete inference cycle, that can be used to characterize workload-specific power distributions that underlie AI inference aggregate power models, providing a foundation for more realistic energy evaluation and infrastructure planning.

BACKGROUND

In 2025, global data centers consumed over 47,000 MW of power, with AI hardware accounting for a growing share, comparable to the national power consumption of France [8]. AI inference power demand is expected to grow rapidly, with a 2022 Google study attributing 60% of its AI energy use to inference [7, 4]. However, existing research mainly focuses on AI training or total inference energy consumption, leaving the aggregate power draw of AI inference poorly understood [1, 3]. Unlike CPU-dominated workloads, modern AI inference relies on high-power GPUs and can cause significant power fluctuations as requests are processed [9]. Therefore, fine-grained power traces and models are needed for grid integration, capacity planning, and energy-saving strategies such as GPU oversubscription and phase splitting [5, 6, 9]. AI inference consists of two key phases: compute-intensive prefill and memory-bound decode [6]. Our measurement framework therefore captures both phases to characterize inference power behavior.

RESEARCH OBJECTIVES

We study how workload and system factors shape AI inference GPU power traces:
- Batching strategy, input/output length, and request type
- Prefill / decode / idle phase decomposition
- GPU type, parallelization method, numerical format
- Aggregation from individual requests to system-level power

AI Inference Workload — Batch Size, Input/Output Length, Request Type
↓
GPU Power Trace Measurement — Prefill / Decode / Idle Phases
↓
Power Characterization — GPU Type, Parallelism, Data Type Effects
↓
Aggregate Power Model — Real-world AI Inference Workloads

Framework overview: request → power trace capture → phase decomposition → aggregate model

EXPERIMENTAL DESIGN

1. Controlled experiments vary input length, output length, and prompt category independently while fixing batch size, model, GPU, and numerical format, isolating the impact of each workload factor on power consumption.
2. Repeat experiments across four GPUs (H100, H200, A100, L40S), 11 model families (1B–72B), five numerical formats (fp16, bf16, fp8, GPTQ, AWQ), and three tensor-parallel settings (TP1, TP2, TP4) to validate workload-power relationships across hardware, scale, and quantization.

⟹ Developed an automated framework to configure experiments and log power traces as .csv files for downstream analysis.

EXAMPLE OF POWER PROFILE

1 • Vary prompt content — input ≈ 5.1-7.7k tok fixed, output 400 tok fixed
novel excerpt | math problems | code completion

2 • Vary input length — novel excerpt content fixed, output 400 tok fixed
short input | medium input | long input

3 • Vary output length — novel excerpt content fixed, input ≈ 8k tok fixed
short output | medium output | long output

KEY FINDINGS

We observed that, at an individual request level, the input and output token lengths were the main factors that affected the GPU power utilization:
- Increasing the input token length caused an increase in the GPU power spikes during the prefill phase.
- Longer output token lengths extended the duration of the decode phase.
- At fixed input and output lengths, the GPU power consumption remained nearly identical across all five prompt types during the decode phase, and showed only small variations for the prefill phase.

⟹ This enables the characterization of workload-specific power distributions simply by considering the collected power traces for fixed input- / output token lengths, that can in turn be used to model the aggregate power.

POWER MODELLING

REQUEST ARRIVAL → GPU RESPONSE → AGGREGATE POWER

Understand the aggregate power of AI inference as a two-part process:
- Request arrivals are modeled as a stochastic process, where each request is described by its arrival time T_a and workload characteristics $Y_a = (l_a^{in}, l_a^{out})$.
- Each active request generates a GPU power response $P_a(t) = f(t - T_a; Y_a)$, sampled from workload-conditioned power distributions obtained from repeated inference measurements.

Combining request arrivals and workload-dependent power responses forms a **Marked Point Process** with mark space $\mathcal{E}$ and attached power response functions [2].
The instantaneous aggregate power is:

$$P_{out}(t) = P_{idle}(t) + \sum_{k \text{ s.t. request is active at time } t} P_k(t)$$

CONCLUSION

With the framework we built it is possible to measure GPU power traces for different AI inference requests, allowing us to study the effects of different factors, especially those affecting the workload such as batch sizes and input- /output token lengths, on the instantaneous power consumption of AI inference requests. We observed that the main factors affecting the GPU power were the input- and output token lengths, which suggests the feasibility of a aggregate power model for AI inference based on power distributions obtained from running AI inference for fixed input and output token lengths.

REFERENCES

[1] E. C. et. al. Power stabilization for ai training datacenters. 2025. URL https://arxiv.org/abs/2508.14318.
[2] M. Jacobsen. *Point Process Theory and Applications - Marked Point and Piecewise Deterministic Processes*. Birkhauser, 2006. ISBN 13 978-0-8176-4215-5.
[3] Jae-Won, Chung, R. Wu, J. J. Ma, and M. Chowdhury. Where do the joules go? diagnosing inference energy consumption. *University of Michigan The ML.ENERGY Initiative*, 2026.
[4] S. Luccioni, Y. Jernite, and E. Strubell. Power hungry processing: Watts driving the cost of ai deployment? In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, page 85–99, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704505. doi: 10.1145/3630106.3658542.
[5] P. Patel, E. Choukse, C. Zhang, I. Goiri, B. Warrier, N. Mahalingam, and R. Bianchini. Characterizing power management opportunities for llms in the cloud. 2024. doi: https://doi.org/10.1145/3620666.3651329.
[6] P. Patel, E. Choukse, C. Zhang, A. Shah, Íñigo Goiri, S. Maleki, and R. Bianchini. Splitwise: Efficient generative llm inference using phase splitting, 2024. URL https://arxiv.org/abs/2311.18677.
[7] D. Patterson, J. Gonzalez, U. Hölzle, Q. Le, C. Liang, L.-M. Munguia, D. Rothchild, D. So, M. Texier, and J. Dean. The carbon footprint of machine learning training will plateau, then shrink, 2022. URL https://arxiv.org/abs/2204.05149.
[8] S. Sujadeh, L. Fattorini, R. Perrault, Y. Gil, V. Parli, L. Santarlasci, J. Pavo, N. Maslej, R. Ahmad, E. Brynjolfsson, C. Bradley, J. Clark, V. Dignam, V. Kumar, J. Landay, T. Lyons, J. Manyika, J. C. Niebles, Y. Shoham, E. Tabassi, R. Wald, T. Walsh, and D. Wald. The ai index 2026 annual report. *AI Index Steering Committee, Institute for Human-Centered AI*, pages 30–35, 2026. doi: https://doi.org/10.48550/arXiv.2606.13206.
[9] G. Wilkins, F. Kazhamiaka, and R. Rajagopal. From servers to sites: Compositional power trace generation of llm inference for infrastructure planning, 2026. doi: https://doi.org/10.1145/nnnnnnn.nnnnnnn.

ACKNOWLEDGEMENTS

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for supporting this project.

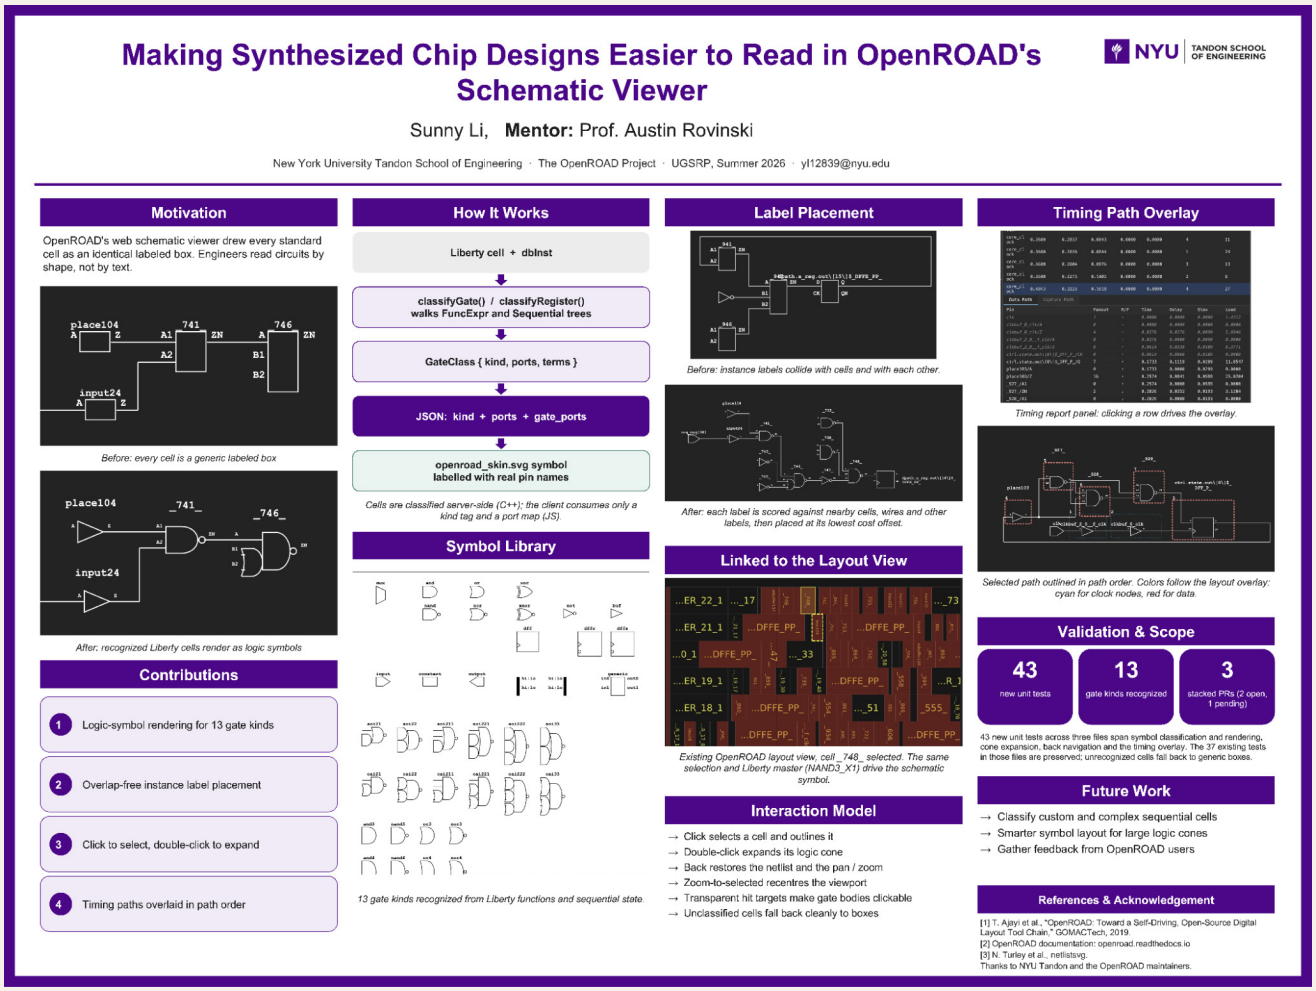
Systems Engineering & Complex Decision-Making



Yucan Li
Electrical Engineering and Computer Engineering (dual major), B.S.
Fall 2028
NYU Tandon
Thompson-Bartlett Fellow

Interactive Logic-Symbol Rendering for OpenROAD Gate-Level Schematics

Engineers working with synthesized digital designs often depend on timing reports to understand how a circuit behaves, but timing alone does not reveal how a design is actually built. Logic gates, how they connect, and how signals flow through the circuit are all important. Without that structural picture, exploring or debugging a design is difficult, much like trying to repair a machine without being able to see inside it. OpenROAD, an open-source RTL-to-layout chip design flow, includes a web-based schematic viewer intended to provide this structural view, but it previously rendered many standard cells as plain, identical boxes, making common logic and sequential elements difficult to recognize at a glance. This project extends the viewer to draw standard cells as familiar logic symbols while preserving OpenROAD's original instance names and pin information. On the backend, schematic generation was extended to classify standard cells such as buffers, inverters, multi-input gates, and flip-flops, then pass the matching symbol information to the frontend, where each cell is drawn as a recognizable shape with its true pin labels. Cells that cannot be identified fall back to the original box view. The result is an interactive viewer that lets users explore a design the way engineers naturally think about it, recognizing gates on sight, tracing signals, and moving through the circuit to focus on the parts that matter. A full set of automated tests confirms the features work reliably. By making synthesized designs visually recognizable, this work gives OpenROAD users structural insight that timing-focused views alone cannot provide, making chip designs easier to explore and understand.



Systems Engineering & Complex Decision-Making



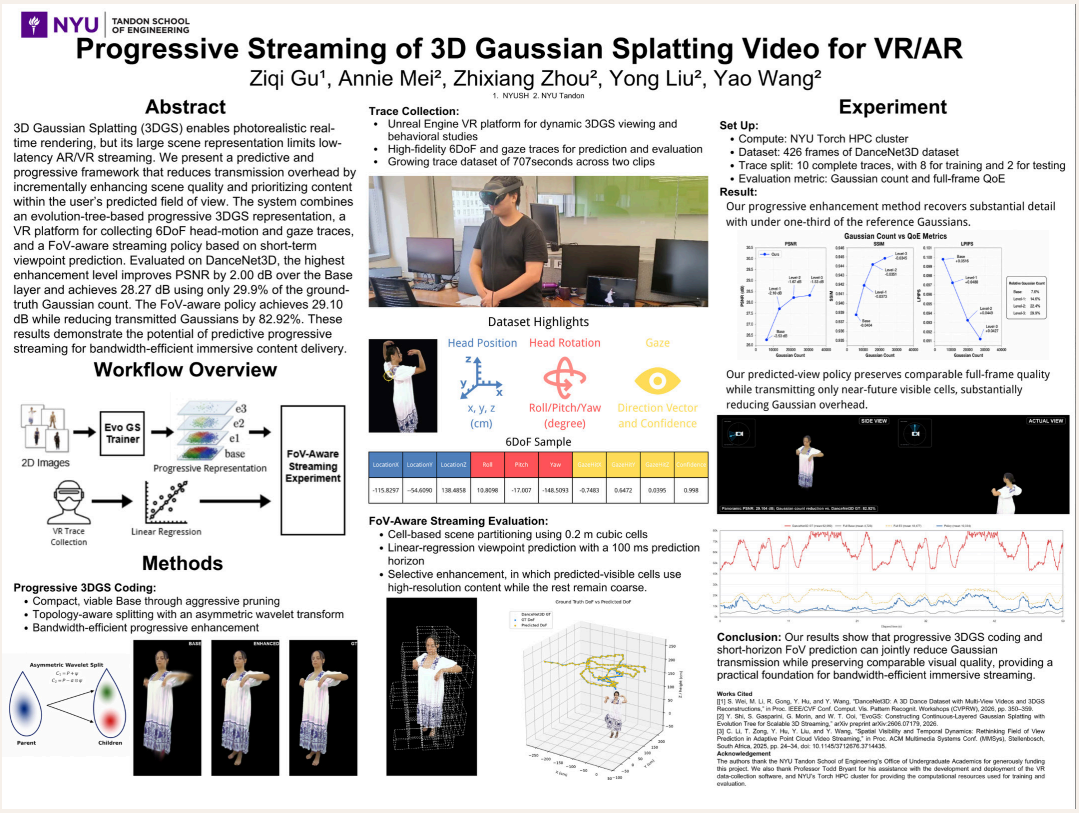
Ziqi Gu
Computer Science, B.S.
Spring 2027
NYU Shanghai



Annie Mei
Computer Engineering, B.S.
Spring 2028
NYU Tandon

Predictive and Progressive Streaming of 3D Gaussian Splatting for AR/VR Applications

3D Gaussian Splatting (3DGS) supports photorealistic, real-time rendering of immersive 3D scenes, but its large data size creates major challenges for low-latency transmission in AR/VR applications. This project explores predictive immersive video streaming that anticipates users' future viewpoints and prioritizes likely visible content, aiming to reduce bandwidth consumption while preserving visual quality. We investigate two complementary approaches. First, we develop a progressive streaming pipeline that generates multiple quality levels of 3DGS content, transmitting a compact base representation first and progressively adding enhancement data to restore visual quality. Second, we use a FoV prediction framework to estimate users' future viewpoints from recent head-motion traces and identify the 3DGS content most likely to be visible. We conduct simulation experiments to evaluate the effectiveness of the proposed design under controlled settings. Preliminary results show that our progressive coding method produces more compact 3DGS representations while achieving higher reconstruction quality than the evaluated baselines. It reduces the number of enhancement Gaussians and overall model size while improving PSNR, indicating a more efficient quality–size trade-off. Initial FoV prediction experiments further suggest that recent user motion provides useful information for estimating future viewing directions and prioritizing potentially visible scene regions. Although prediction accuracy varies with the prediction horizon and user movement patterns, the results support the feasibility of incorporating viewpoint prediction into adaptive 3DGS streaming. These results demonstrate the potential of predictive and progressive streaming to make immersive 3D content more practical for bandwidth-constrained AR/VR applications. By reducing transmission costs while maintaining visual quality, the proposed framework provides a promising foundation for scalable and low-latency immersive video delivery.





Engineering & Culture

Engineering and culture are not separate domains, but two sides of the same coin. We are cultivating a new generation of technologists who bring together left brain and right brain thinking to solve complex problems and create human-centered experiences. Our research and programs blend technical rigor with creativity to design solutions and environments that preserve knowledge, convey knowledge, open minds, and inspire joy.

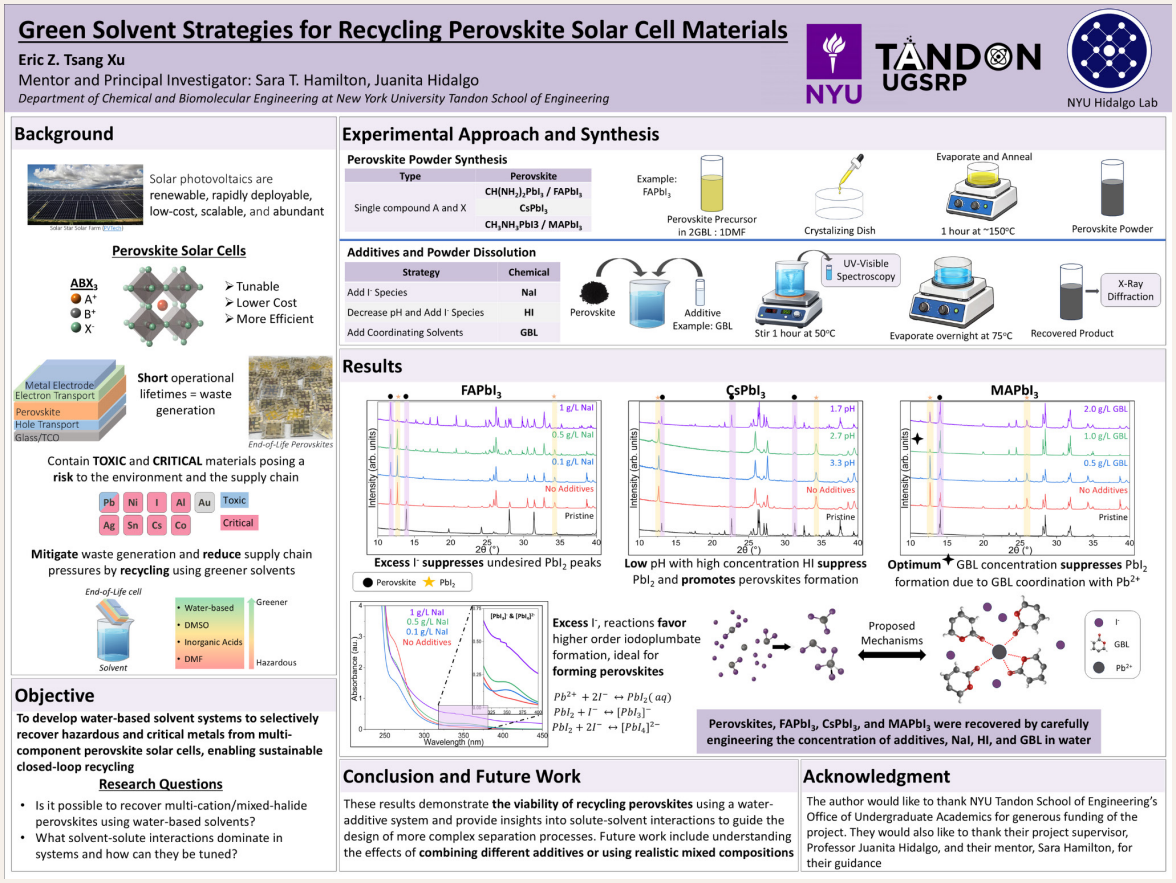


Eric Tsang Xu
Chemical and Biomolecular Engineering, B.S.
Spring 2028
NYU Tandon

Green Solvent Strategies for Recycling Perovskite Solar Cell Materials

Perovskite solar cells (PSCs) have attracted significant research attention owing to their exceptional performance and low cost. However, their short operational lifetimes related to environmental and thermal instabilities, and inclusion of hazardous and critical materials (e.g., lead (Pb) and cesium (Cs)) make it essential to develop sustainable recycling techniques for PSCs to mitigate hazardous waste generation and reduce supply chain pressures. Current research on recovering end-of-life perovskites has yet to report the systematic recovery of multi-cation/mixed-halide perovskites using green solvents like water. Additionally, the lack of fundamental understanding of solvent-solute interactions hinders rational design of perovskite recycling strategies.

In this work, we investigated lead halide perovskite recovery across a compositional library using water and simple additives. We systematically varied the A-site cation of perovskite powder APbI3 (A = MA, FA, Cs) synthesized by evaporating and annealing a precursor solution. To understand composition and additive dependent dissolution and recovery behavior, the perovskite powders were dissolved in water with NaI, gamma-butyrolactone (GBL), or HI additives. Additional I- is hypothesized to promote iodoplumbate complex formation, GBL to promote PbI2 dissolution, and HI to decrease pH and increase solute solubility. To gain insights on the solvent-solute interactions in perovskite solution with and without additives, UV-Visible spectroscopy was employed. Absorption spectra indicated the formation of higher order iodoplumbate complexes with I- addition or optimal GBL concentration. X-ray diffraction was used to determine the phase purity and crystallinity of the products. Powders recovered without additives consisted mostly of PbI2 and other impurities. In contrast, by controlling additive concentration and solution pH, perovskites could be recovered from aqueous solutions to varying extents depending on composition. While future work is needed, these initial results demonstrate the viability of recycling perovskites using a water-additive system and provide insights into solute-solvent interactions to guide the design of more complex separation processes.





Genesis Tamara Reyes Casas
Interactive Media , B.S.
Spring 2027
NYU Abu Dhabi

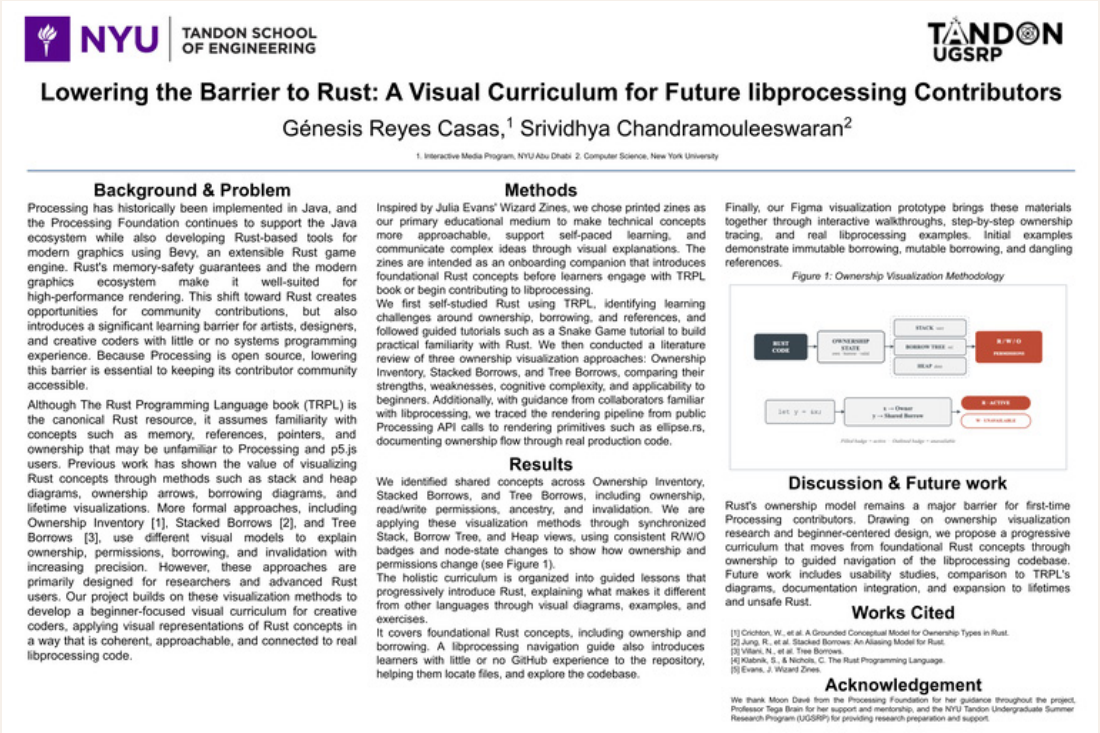
Lowering the Barrier to Rust: A Visual Curriculum for Future libprocessing Contributors

The Processing Foundation's historically Java-based tooling is now expanding to Rust, creating an opportunity to modernize and extend its open-source ecosystem while introducing a significant learning barrier for artists and creative technologists with little or no systems programming experience. This project investigates how beginner-friendly educational materials can help creatives develop foundational Rust knowledge before engaging with libprocessing as contributors.

Our approach combines a review of existing research on memory-safety education with the design of a visual, example-driven curriculum. Drawing from The Rust Programming Language, Rust by Example, and p5.js learning resources, we examine how concepts familiar to creative coders can serve as entry points into unfamiliar systems programming ideas. Research on Rust education suggests that beginners often recognize compiler errors without understanding the underlying concepts, such as why a dangling reference is unsafe. This gap between identifying an error and understanding its cause motivates our focus on visual explanations that make abstract concepts such as ownership, borrowing, and permissions easier to understand.

Rather than beginning with detailed representations of Rust's internal memory model, the curriculum introduces concepts incrementally through concrete examples and visualizations of how data can be accessed, moved, and modified throughout a program. We are synthesizing insights from three models of Rust memory safety—Permissions, Stacked Borrows, and Tree Borrows—to develop a visualization hierarchy that communicates these relationships without overwhelming beginners. The curriculum is presented as a series of printable zines supported by Figma templates, allowing the materials to evolve as the project and community grow.

By connecting Rust concepts to the visual and interactive learning practices familiar to creative coders, this project aims to bridge the gap between creative coding and systems programming. We expect the resulting curriculum to provide an accessible starting point for learners who may be intimidated by traditional systems programming education, while establishing a foundation for future contributors of libprocessing.



NYU TANDON SCHOOL OF ENGINEERING

TANDON UGSRP

Lowering the Barrier to Rust: A Visual Curriculum for Future libprocessing Contributors

Génesis Reyes Casas,[1] Srividhya Chandramouleeswaran[2]

1. Interactive Media Program, NYU Abu Dhabi 2. Computer Science, New York University

Background & Problem

Processing has historically been implemented in Java, and the Processing Foundation continues to support the Java ecosystem while also developing Rust-based tools for modern graphics using Bevy, an extensible Rust game engine. Rust's memory-safety guarantees and the modern graphics ecosystem make it well-suited for high-performance rendering. This shift toward Rust creates opportunities for community contributions, but also introduces a significant learning barrier for artists, designers, and creative coders with little or no systems programming experience. Because Processing is open source, lowering this barrier is essential to keeping its contributor community accessible.

Although The Rust Programming Language book (TRPL) is the canonical Rust resource, it assumes familiarity with concepts such as memory, references, pointers, and ownership that may be unfamiliar to Processing and p5.js users. Previous work has shown the value of visualizing Rust concepts through methods such as stack and heap diagrams, ownership arrows, borrowing diagrams, and lifetime visualizations. More formal approaches, including Ownership Inventory [1], Stacked Borrows [2], and Tree Borrows [3], use different visual models to explain ownership, permissions, borrowing, and invalidation with increasing precision. However, these approaches are primarily designed for researchers and advanced Rust users. Our project builds on these visualization methods to develop a beginner-focused visual curriculum for creative coders, applying visual representations of Rust concepts in a way that is coherent, approachable, and connected to real libprocessing code.

Methods

Inspired by Julia Evans' Wizard Zines, we chose printed zines as our primary educational medium to make technical concepts more approachable, support self-paced learning, and communicate complex ideas through visual explanations. The zines are intended as an onboarding companion that introduces foundational Rust concepts before learners engage with TRPL book or begin contributing to libprocessing.

We first self-studied Rust using TRPL, identifying learning challenges around ownership, borrowing, and references, and followed guided tutorials such as a Snake Game tutorial to build practical familiarity with Rust. We then conducted a literature review of three ownership visualization approaches: Ownership Inventory, Stacked Borrows, and Tree Borrows, comparing their strengths, weaknesses, cognitive complexity, and applicability to beginners. Additionally, with guidance from collaborators familiar with libprocessing, we traced the rendering pipeline from public Processing API calls to rendering primitives such as ellipse.rs, documenting ownership flow through real production code.

Results

We identified shared concepts across Ownership Inventory, Stacked Borrows, and Tree Borrows, including ownership, read/write permissions, ancestry, and invalidation. We are applying these visualization methods through synchronized Stack, Borrow Tree, and Heap views, using consistent R/W/O badges and node-state changes to show how ownership and permissions change (see Figure 1).

The holistic curriculum is organized into guided lessons that progressively introduce Rust, explaining what makes it different from other languages through visual diagrams, examples, and exercises.

It covers foundational Rust concepts, including ownership and borrowing. A libprocessing navigation guide also introduces learners with little or no GitHub experience to the repository, helping them locate files, and explore the codebase.

Finally, our Figma visualization prototype brings these materials together through interactive walkthroughs, step-by-step ownership tracing, and real libprocessing examples. Initial examples demonstrate immutable borrowing, mutable borrowing, and dangling references.

Figure 1: Ownership Visualization Methodology

Discussion & Future work

Rust's ownership model remains a major barrier for first-time Processing contributors. Drawing on ownership visualization research and beginner-centered design, we propose a progressive curriculum that moves from foundational Rust concepts through ownership to guided navigation of the libprocessing codebase. Future work includes usability studies, comparison to TRPL's diagrams, documentation integration, and expansion to lifetimes and unsafe Rust.

Works Cited

[1] Crichton, W., et al. A Grounded Conceptual Model for Ownership Types in Rust.
[2] Jung, R., et al. Stacked Borrows: An Aliasing Model for Rust.
[3] Villani, N., et al. Tree Borrows.
[4] Klabnik, S., & Nichols, C. The Rust Programming Language.
[5] Evans, J. Wizard Zines.

Acknowledgement

We thank Moon Davé from the Processing Foundation for her guidance throughout the project, Professor Tega Brain for her support and mentorship, and the NYU Tandon Undergraduate Summer Research Program (UGSRP) for providing research preparation and support.



Kacper Jucha
Mechanical Engineering, B.S.
Spring 2027
NYU Tandon

Integrating Artificial Intelligence into First Year Engineering Coursework and Evaluating its Impact on Technical Skill Development

Artificial Intelligence (AI) is becoming integrated into many aspects of our lives, from making to-do lists, to analyzing data. The effective use of AI is a growing expectation of engineering graduates; still, AI tools have not been widely adopted in the classroom. This research aims to integrate AI in a first year engineering design course, EG-1004 Introduction to Engineering and Design. Several tools were piloted to help students use AI as a learning partner in their work, rather than a tool which completes their work for them. Initial analysis began with a Pre-Test to determine the engineering and professional skills students entered the classroom with. Generally, students were found to enter the classroom with good professional engineering skills of design, communication, data analysis, yet were lacking technical engineering skills such as CAD, coding, and circuitry. Experience with AI tools varied among the students. In response, customized Google AI assistants, known as Gemini Gems, were developed in order to aid as peer tutors in the students' laboratory environment. Part of this included researching the limitations of the software and identifying the point at which teaching assistant intervention is required. This preliminary research is guiding a larger rollout for the 2026-27 academic year, where a Post-Test will track the success of these AI tools in the development and retention of technical engineering skills. These tests will establish a standardized framework for determining the efficacy of the coursework and provide insightful data for future revisions of the course curriculum. Students must be able to use AI to their advantage and by setting a standard of balanced AI practice in the first year, this research proposes that they will be able to utilize AI more effectively for their future classes and careers because of these early academic interventions.

NYU TANDON SCHOOL OF ENGINEERING

GENERAL ENGINEERING | TANDON UGSRP

Integrating Artificial Intelligence into First Year Engineering Coursework and Evaluating its Impact on Technical Skill Development

Kacper Jucha[1], Ingrid Paredes[2]
1. R&D Lead, General Engineering, 2. Director, General Engineering

Abstract

Artificial Intelligence (AI) is becoming integrated into many aspects of our lives, from making to-do lists, to analyzing data. The effective use of AI is a growing expectation of engineering graduates; still, AI tools have not been widely adopted in the classroom. This research aims to integrate AI in a first-year engineering design course, EG-1004 Introduction to Engineering and Design. Several tools were piloted to help students use AI as a learning partner in their work, rather than a tool which completes their work for them. Initial analysis began with a Pre-Test to determine the engineering and professional skills students entered the classroom with. Generally, students were found to enter the classroom with good professional engineering skills of design, communication, data analysis, yet were lacking technical engineering skills such as CAD, coding, and circuitry. Experience with AI tools varied among the students. In response, customized Google AI assistants, known as Gemini Gems, were developed in order to aid as peer tutors in the students' laboratory environment. Part of this included researching the limitations of the software and identifying the point at which teaching assistant intervention is required. This preliminary research is guiding a larger rollout for the 2026-27 academic year, where a Post-Test will track the success of these AI tools in the development and retention of technical engineering skills. These tests will establish a standardized framework for determining the efficacy of the coursework and provide insightful data for future revisions of the course curriculum. Students must be able to use AI to their advantage and by setting a standard of balanced AI practice in the first year, this research proposes that they will be able to utilize AI more effectively for their future classes and careers because of these early academic interventions.

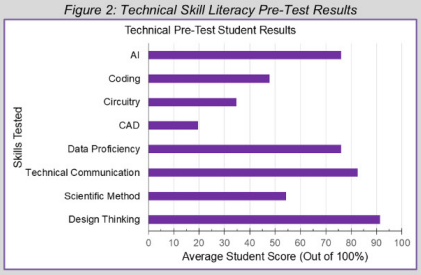
Testing

- Pre-Test and standardized rubric developed to ensure reproducibility of testing (Technical skill testing typically lacks reproducibility and scale)
- Example Assessment Question: *"Describe the steps of the scientific method in your own words. You can write this in list or paragraph form."*
- Testing was done in a written paper format in a standard classroom setting where students were not allowed to collaborate or use technology
- Students were not informed whether the Pre-Test (ungraded) had an impact on their grade as to preserve their attention on the assessment
- Students were informally provided with trained Gemini Gems alongside their laboratory coursework to engage with and provide feedback

Figure 1: Sample of Rubric for Pre-Test

Skill/Section	0 Point	1 Points	2 Points
(A) Design Thinking	Did not attempt or did not answer the provided question.	Attempted the sketch and the sketch is somewhat explained. Overall average performance (some labels included, sloppy) or did not meet all written requirements for the sketch.	Sketch is neatly drawn, parts are labeled, and the sketch effectively communicates all provided requirements. Description of sketch is concise, clear, and communicates what is in the sketch.
(B) Scientific Method	Did not attempt or did not answer the provided question.	Does not list all steps of the scientific method or explanation of the scientific method lacks detail. Overall average answer which somewhat communicates what the scientific method is.	Lists all steps of the scientific method or provides a clear, concise, and correct explanation of the scientific method.
(C) Technical Communication	Did not attempt or did not answer the provided question.	Included the types of technical writing they have done and/or attempted to explain a concept of their choosing. The explanation poorly communicates the concept or does not provide the detail which an educated, but uninformed individual would require to understand the concept.	Includes the types of technical writing they have done and explains a concept of their choosing. The explanation is concise, includes necessary requirements, and explains the technical concept in a detailed manner.

Figure 2: Technical Skill Literacy Pre-Test Results



Results & Conclusions

- Pre-Test: students enter their first year of engineering with greater professional engineering skills, lack technical engineering skills such as CAD, circuitry, and coding (<50% average student score)
- AI Skills: Students showed stronger, yet varied prompting skills
- Informal feedback about Gemini Gems: Students appreciated the guiding tone of the bot, along with helping answer simple questions and referring to teaching assistants when more difficult questions arose

Integration of AI into Curriculum

- Gemini Gems were trained on internal documents and given strict prompts to remain within academic boundaries
- Coding Lab: Gemini Gem "peer tutor" functions as an assistant for students and guides them along any questions they may have without providing the solution
- Ideation Lab: Gemini Gem "consultation bot" functions as a random individual with a problem and asks guiding questions for students to think critically about approaching/solving a real-world problem
- Project Workshop: As students research their project, they will generate images of their ideas using Gemini and discover the benefits and shortcomings of generative AI

Future Work

- Development of a Post-Test to determine coursework effectiveness on technical skill development and the efficacy of implemented tools such as AI
- Implementation of AI into more sectors of the curriculum
- Use Pre/Post-Test data to quide curriculum development

Works Cited

[1] Garg, Ashish et al. "Enhancing data analysis and programming skills through structured prompt training: The impact of generative AI in engineering education." Comput. Educ. Artif. Intell. 8 (2025): 100380.
[2] Mehek Kunal Vora and Ethan E Danahy. "A Student Classification and Characterization Model of Generative AI Use in First-Year Engineering Design". 2025 ASEE Annual Conference & Exposition , Montreal, Quebec, Canada , 2025. June ASEE Conferences, 2025. https://peer.asee.org/55398.internet.
[3] Alsharefeen, Rami, and Naji Al Sayari. "Examining Academic Integrity Policy and practice in the era of AI: A case study of faculty perspectives." Frontiers in Education, vol. 10, 26 June 2025, https://doi.org/10.3389/feduc.2025.1621743.

Acknowledgement

The authors thank NYU Tandon School of Engineering's Office of Undergraduate Academics for generous funding of the project.

Engineering & Culture



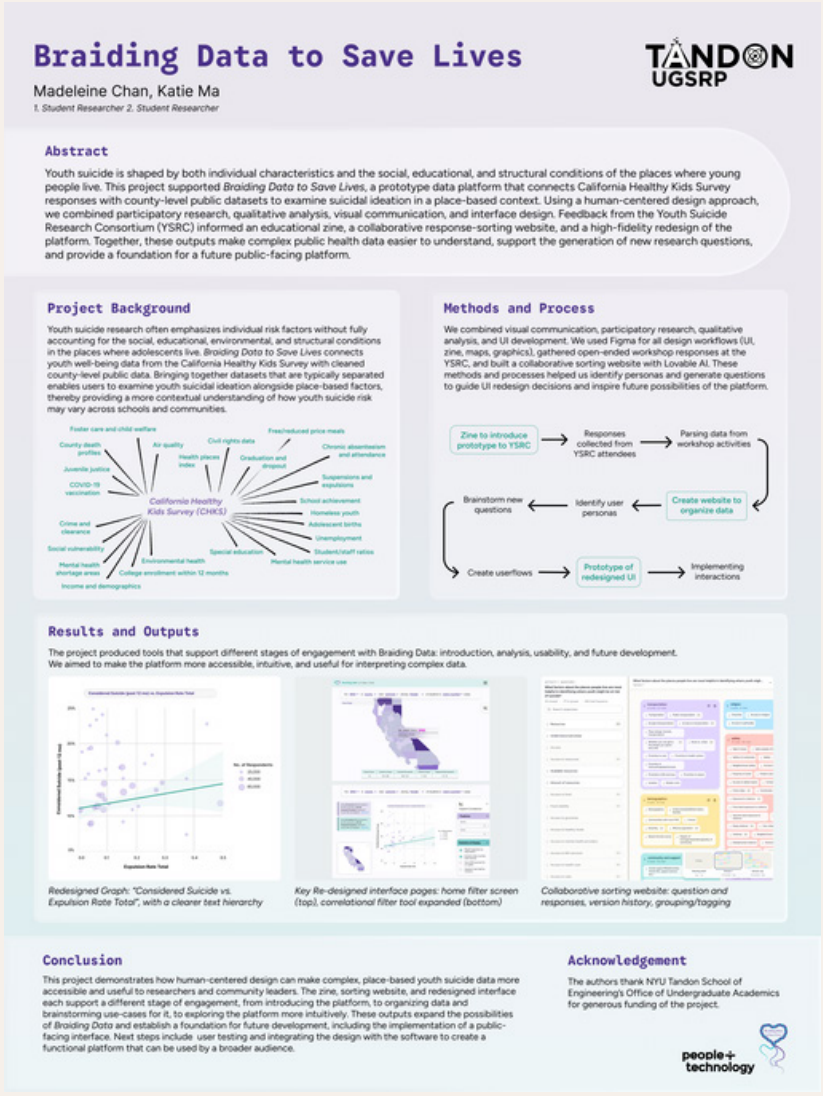
Madeleine Chan
Integrated Design & Media, B.S.
Spring 2027
NYU Tandon

Braiding Data to Save Lives

Youth suicide research often emphasizes individual risk factors without fully accounting for the social, educational, and structural conditions in the places where adolescents live. Braiding Data to Save Lives is a prototype open R data system that links student well-being responses from the California Healthy Kids Survey with cleaned county-level public data from education, health, census, environmental, and justice agencies, enabling users to examine youth suicidal ideation in a place-based context.

Our work combined visual communication, participatory research, qualitative analysis, and interface design. We created a zine to introduce the platform and its potential uses to attendees of the Youth Suicide Research Consortium, then analyzed workshop participants' written responses to identify perspectives, needs, and possible applications. To support collaborative synthesis, we developed a custom website through which team members could flexibly group responses, identify user personas and motivations, and generate questions and insights.

We also used a user-centered design process to improve the platform's interface and data visualizations. Researcher personas and user workflows guided the organization of information, while user feedback informed iterative refinements. We produced a high-fidelity Figma prototype that streamlined navigation and improved how users interact with data, graphs, tables, filters, and maps. The redesigned interface supports clearer exploration of relationships across demographic, geographic, and socioeconomic factors and creates a foundation for future analytical and predictive tools.





Mark Ionis
Chemical and Biomolecular Engineering, B.S.
Spring 2029
NYU Tandon

Geometric Deep Learning for Structure-Based Enzyme Kinetic Prediction

Enzyme turnover numbers (kcat), Michaelis-Menten constants (Km), and catalytic efficiencies (kcat/Km) are key kinetic parameters used to quantify catalytic efficiency for enzyme engineering applications. Experimental methods for determining these parameters are expensive and time-intensive, leading to the development of machine learning-based predictive models. However, many of those models only utilize enzyme sequences and substrate representations as inputs, limiting their ability to accurately capture enzyme-substrate interactions. To address these limitations, we developed a geometric deep learning model to predict kcat, Km, and kcat/Km for β-glucosidases, a key enzyme family used in biorefineries, used here as a case study. Molecular docking was used to represent the position of the ligand within the active site of the enzyme, and the resulting complex was converted to a graphical network. Preliminary results suggest that this structure-informed model performed comparably to a sequence-only baseline. The structure-informed model performs much better than the baseline at predicting Km with a mutant-blind split, however. This suggests that a structure-based model is better able to generalize to unseen enzymes. Ongoing work involves sampling alternate conformations of the enzyme using molecular dynamics to provide a comprehensive view of its interactions with the substrate. This approach is an endeavor to shift toward physics-informed machine learning that goes beyond statistical patterns in the dataset.

NYU TANDON UGSRP

Geometric Deep Learning for Kinetic Parameter Prediction in β-Glucosidases

Mark Ionis, Ali Malli, Lefteris Komvoupolos, Jin Ryoun Kim
Department of Chemical and Biomolecular Engineering, NYU Tandon School of Engineering, Brooklyn, NY USA

INTRODUCTION

- β-glucosidases (BGLs) are enzymes that hydrolyze cellobiose into glucose and are essential for biorefineries and circular economy
- Machine learning (ML) models can be used to discover more efficient BGLs with higher catalytic efficiencies (k_{cat}/K_m)
- Existing ML models trained on global datasets have poor performance on local datasets
- Many of those models only leverage enzyme sequence information without accounting for structural interactions with the substrate

OBJECTIVE & METHODS

- **Objective:** Leveraging structural information from enzyme-substrate complexes for K_m, k_{cat}, and k_{cat}/K_m prediction
- **Scope:** BGLs acting on p-nitrophenyl glucopyranoside (pNP-Glc), a synthetic substrate akin to cellobiose, are used as a case study

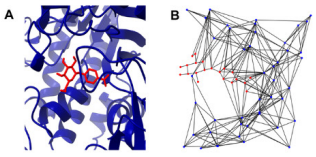


Fig. 1. (A) A BGL-pNP-Glc active site and (B) a projection of its corresponding network.

- **Data Preparation:** Enzyme sequences encoded using ESM2; ligand representations encoded using ChemBERTa-3
- **Graph Construction:** Active site pocket represented as a network, each node is populated with an enzyme or ligand embedding
- **Model Architecture:** 3-layer graph convolutional network (SAGE), attention aggregate pooling, 2-layer multi-layer perceptron (MLP)

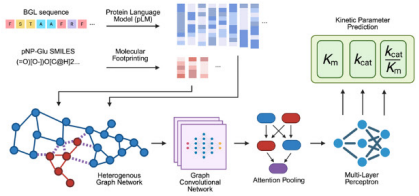


Fig. 2. Graph encoder architecture and inference workflow of the developed model.

RESULTS

- Trained the model on three tasks to predict (i) K_m, (ii) k_{cat}, and (iii) k_{cat}/K_m
- Prediction accuracy was highest for K_m, with k_{cat}/K_m performing similarly

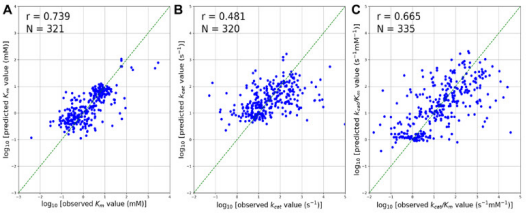


Fig. 3. Comparison between predicted and observed values for (A) K_m, (B) k_{cat}, and (C) k_{cat}/K_m for published enzyme-substrate pairs of BGLs and pNP-Glc, acquired using a random split. Pearson Correlation Coefficients (r) and sample sizes (N) are shown.

- There is a data leakage threat as 51% of the dataset is mutant sequences with high similarity to their wild-type counterparts
- Evaluating the model on a 99% sequence-identity split limited prediction accuracy, but retained a signal

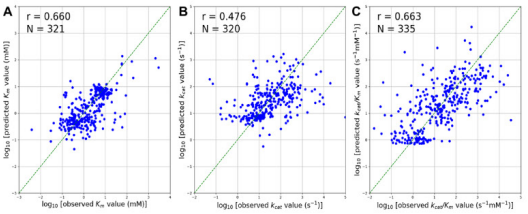


Fig. 4. Comparison between predicted and observed values for (A) K_m, (B) k_{cat}, and (C) k_{cat}/K_m for published enzyme-substrate pairs of BGLs and pNP-Glc, acquired using a 99% sequence identity split. Pearson Correlation Coefficients (r) and sample sizes (N) are shown.

- Consistent with the literature, the model performs more accurately on mutant sequences than on wild-type sequences due to their similarity

Table 1. Pearson correlational coefficients (r) and Spearman's ρ for models trained on wild-type (WT) and mutant sequences for all kinetic parameters.

Enzyme	K_m		k_{cat}		k_{cat}/K_m	
	Mean r	Mean ρ	Mean r	Mean ρ	Mean r	Mean ρ
WT	0.561	0.474	0.272	0.345	0.405	0.409
Mutant	0.809	0.74	0.617	0.607	0.79	0.763

- The structure-based model performed comparably to a sequence-only baseline
- Models remain comparable regardless of split

Table 2. Pearson correlational coefficients (r) and Spearman's ρ for the structure-based model and an MLP baseline for all kinetic parameters.

Model	Metric	Random Split		99% Similarity Split	
		Mean r	Mean ρ	Mean r	Mean ρ
Structure-Based	K_m	0.739	0.724	0.66	0.546
	k_{cat}	0.481	0.538	0.476	0.486
	k_{cat}/K_m	0.665	0.68	0.633	0.604
Sequence-Based	K_m	0.724	0.685	0.667	0.538
	k_{cat}	0.469	0.533	0.458	0.447
	k_{cat}/K_m	0.625	0.631	0.664	0.636

CONCLUSION

- The structure-based model successfully predicted K_m, k_{cat}, and k_{cat}/K_m for BGLs comparably to a model trained only on sequence
- Structural data provides more accurate prediction for K_m than k_{cat}
- Highly similar mutants in the dataset limit model generalizability

NEXT STEPS

- Incorporate features from molecular dynamics (MD) to better represent enzyme-substrate interactions
- Use transfer learning to add K_m and k_{cat} weights to k_{cat}/K_m prediction
- Decrease noise in predictions by reducing dimensions with PCA

REFERENCES

1. Ali Malli, Denys Vasyutyn, Jin Ryoun Kim; Advances in Machine Learning Models for Predicting Enzyme Kinetic Parameters. *J. Chem. Inf. Model.* 12 Jan 2026; 66(1): 42-60.
2. Ding Luo, Huining Ji, Shuming Cheng, Kaiqi Wen, Xiaoyang Qu, Mingfeng Cao, Liang Hong, Binju Wang; Enzyme Kinetic Parameter Prediction via Catalytic Pocket-Augmented Machine Learning. *ACS Catal.* 3 July 2026; 16 (13): 12784–12801.
3. Mehmet Emre Erkanli, Yunseok Jang, Ali Malli, Khalid El-Halabi, Chaehyun Ryu, Jin Ryoun Kim; A Three-Module Machine Learning Framework for Protein Sequence- and Temperature-Dependent k_{cat}/K_m Prediction in β-Glucosidases. *ACS Synth. Biol.* 17 October 2025; 14 (10): 3927–3939.

ACKNOWLEDGMENTS

The authors would like to thank NYU Tandon School of Engineering's Undergraduate Summer Research Program for their generous funding of their project. Special thanks to Ali Malli and Professor Jin Ryoun Kim for the opportunity to participate in this project.

Humboldt Internship Program (HIP)



SANTIAGO LACAYO

This past summer, I interned at Doings AI, an AI startup in Berlin, where I worked on front-end development for their app as part of the HIP program. It was such a great experience overall, and the two things I loved the most were, first, learning how an AI tech startup operates from the inside, and second, being part of the broader startup world. The office was shared with many other startups, which I found really cool since I got to meet people and learn about all kinds of creative ideas being built around me. The exchange also gave me the chance to experience Berlin's culture, and I have to say I was amazed by how beautiful the city is and how much there is to do, from its history and unique architecture, mostly Bauhaus and brutalist, to museums like the Neue Nationalgalerie, the people, and even the techno nightlife, which in Berlin is genuinely considered part of the culture.





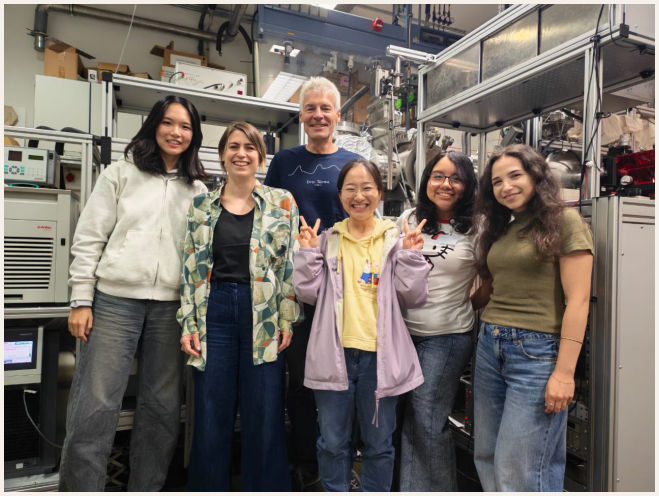
This whole experience gave me many memorable moments and changed me a lot, both professionally and personally. Working alongside people from a different country and exchanging knowledge in that environment taught me how to adapt quickly and collaborate across cultures, skills I am eager to keep building as I continue my studies.

Humbolt Internship Program (HIP)

GARVV GOEL

The Humboldt Internship Program gave me a wonderful experience over the summer, getting to spend real time in Berlin. Humboldt has such a beautiful campus, surrounded by statues and beautiful architecture. It was definitely weird at first getting used to new cultural norms, like not jaywalking or navigating the extensive public transit. The program brought together an incredible mix of people from all over the world, and gave us countless ways to explore the city together, from beach days, beer gardens, billiards nights, and many more ways to experience Berlin more as a resident and less as a tourist. The program also provided trips outside the city to places like Potsdam. Professionally, the internship offered a real look into German work culture, and presenting at the symposium alongside everyone else's projects was a highlight. I personally lived in Mehringdamm, in the heart of Kreuzberg, which meant embracing a neighborhood that was lively, sometimes loud with sirens, but full of energy, especially during World Cup watch parties. The area was crowded with local Spätis and restaurants, making it so convenient to grab food. Beyond daily life, I loved diving into Berlin's history, visiting Berlin Wall sites, Checkpoint Charlie, and various museums. Altogether, HIP was a true cultural immersion, one that let me meet people from around the world and really experience Berlin.

SIYU MA



Spending ten weeks in Berlin for my internship at Humboldt University was an incredible experience! I worked under Professor Wolfgang Christen conducting research in physical chemistry about propane. I learned so much about both physical procedures and analytical skills in the lab. Wolfgang is a wonderful professor alway patient, optimistic and funny. Since it was my very first time visiting Europe, I chose to live in the NYU Berlin dorms. It turned out to be a marvelous choice, and the dorm environment made me feel right at home.

We went on great trips together, like to Potsdam, and I spent plenty of time exploring Berlin itself. I visited Museum Island with its world-class museums to see the famous Bust of Nefertiti, and I even got to experience a pre-season friendly match with Union Berlin. I completely fell in love with the city's unique vibe!

I made wonderful friends through the HIP internship from all around the world—from Singapore, Chile, and France to Kazakhstan. It was an amazing mix of cultures, and I'm so happy to have connected with everyone. The energy in Germany was unforgettable, especially during the World Cup. Even though Germany was knocked out early, the people's passion for football never faded, and I ended up watching the final with a group of strangers, sharing snacks and cheering together like old friends. I also got the chance to travel around the country. I visited Munich, where it was beautiful and sunny near the Alps. I was incredibly lucky to be there right when the local brewers graduated (a special event that only happens once every two years) and enjoyed some wonderful, fresh beer. Another huge highlight was traveling to Hamburg to see Miniatur Wunderland. It felt so unreal yet incredibly detailed and alive, with tiny displays that actually moved and lit up. Looking back, I feel so grateful for this entire opportunity, returning home with great research experience and memories I will keep forever.

KASHVI GUPTA

During my Host Internship Program experience in Berlin, I interned at arcneo, a FinTech startup, where I worked on lead research for VuV and Stiftungen (foundations) and created AI marketing automation. I also had the opportunity to work closely with the founding team, which gave me a firsthand look at how decisions get made in a fast-moving startup environment and how much ownership you can take on even as an intern. Beyond the professional experience, living and working in Berlin pushed me to immerse myself in a new culture, meet people from completely different backgrounds, and take on challenges I hadn't faced before. The experience sharpened both my technical and interpersonal skills and gave me a much broader perspective on how startups operate outside the U.S. It was an incredibly rewarding few months, both personally and professionally.



Lab Link-Up Summer Social and Luna Park













