

About the Symposium

This symposium presents the work of undergraduate students who spent summer 2026 studying questions in pure mathematics, mathematical modeling, simulation, data science, and related interdisciplinary areas.

The projects in this brochure were completed through two programs in the Department of Mathematics, CMUSR and CSUREMM. Each program ran for ten weeks, giving students the opportunity to work closely with faculty, graduate students, and other mentors on sustained undergraduate research projects.

CMUSR: Columbia Math Undergraduate Summer Research

CMUSR gives students the opportunity to work in small groups on research problems in pure and applied mathematics. Each group works with a faculty adviser and a graduate student mentor, with an emphasis on learning the relevant mathematics, developing new ideas, and working collaboratively toward original results. CMUSR is supported by the Department of Mathematics.

Faculty Leads: Sven Hirsch and Tommaso Maria Botta
Graduate Mentors: Raphael Tsiamis and Igor Chaban

CSUREMM: Columbia Summer Undergraduate Research Experiences in Mathematical Modeling

CSUREMM brings students together in teams to study open-ended problems through mathematical and computational modeling. Students formulate research questions, work with data, develop and test models, and refine their ideas as the projects evolve. CSUREMM receives support from the Department of Mathematics, the Department of Statistics, the Irving Institute for Cancer Dynamics, and the Barnard Summer Research Institute.

Program Lead: George Dragomir

Program Mentors: Beiji Chen (PhD student, APAM) and Andrea Clark (PhD student, Computer Science)

Additional Guidance: Dobrin Marchev (Statistics) and Jiawei Lou (former CSUREMM participant)

We gratefully acknowledge partial support from
Jane Street for Columbia's summer
undergraduate programs.



2026

Undergraduate Summer Research Symposium



Department of Mathematics

Columbia University

Mass, Corners, and Shields

Carmen Alves Sabin, Andrew Arloro, Boqian Qin, Ela Wizel

Faculty Adviser: Sven Hirsch

Graduate Mentor: Raphael Tsiamis

Partially supported by a grant from Simon Brendle.

We prove a Riemannian positive mass theorem for asymptotically flat spin manifolds with corners. In contrast with earlier results, our theorem allows both mean-convex and mean-concave corner regions, subject to a balancing condition. Our strategy expands on the lock principle for scalar curvature introduced by Frenck, Hanke, and Hirsch, as well as the dominant energy shields introduced by Cecchini, Lesourd, and Zeidler.

Duality Interfaces in 3D Mirror Symmetry

Yousif Imad Elhag, David Guikema, Cassidy Kao, Tianyi Lu, Dmytro Prosiannikov, Kathy Xu

Faculty Adviser: Tommaso Maria Botta

Graduate Mentor: Igor Chaban

Partially supported by a grant from Andrei Okounkov.

Symplectic duality exchanges the Higgs and Coulomb branches of 3d, $\mathcal{N} = 4$ supersymmetric gauge theories, allowing geometric information from one branch to be translated into information about its mirror. This correspondence is governed by a duality interface, a section of a line bundle whose fixed-point restrictions are described by elliptic stable envelopes. Although such interfaces are known to exist for certain varieties, existing proofs do not give an explicit construction. In this project, we study the case of the cotangent bundle of a Grassmannian, $T^*\text{Gr}(k, n)$, with the goal of making the duality interface concrete. We reduce the relevant conjecture of Aganagic and Okounkov to finding a quadratic form satisfying certain conditions, propose an explicit candidate formula, and verify the required identities for $k, n < 11$. These results provide strong evidence for the conjecture in the Grassmannian case and identify a candidate line bundle giving the duality interface.

Uncovering Our Galaxy’s Darkest Secret: Improving Black Hole Video Reconstruction with Optimal Transport Regularization

Yefan (Emma) Jia, Ethan Sharma, Raymond Zhao

Sagittarius A*, the black hole at the center of our galaxy, evolves on intrahour timescales, making video reconstruction from Event Horizon Telescope data highly under-constrained. To improve dynamic reconstructions from sparse and noisy observations, we introduce an optimal transport-based temporal regularization method designed to encourage physically consistent motion of brightness across frames while allowing for local brightness changes, such as those caused by flares. We test this approach on simulated black hole data generated from physically motivated models and sampled with realistic Event Horizon Telescope sparsity and noise. Compared with existing static and dynamic reconstruction methods, our signed residual unbalanced optimal transport regularization improves structural similarity by up to 5% and spatiotemporal gradient error by up to 20%, while maintaining comparable chi-squared data fit. These results suggest that optimal transport regularization is a promising approach for improving black hole video reconstruction from very sparse data.

Think Before You Trust: Structure-Aware Verification Budget Allocation in LLM Reasoning

Xinyun (Alina) Du, Christopher Lee, Allison Li

Large language models generate reasoning as networks of dependent claims, where early errors can contaminate downstream conclusions. Since verifying every step is often impractical, we study how a limited verification budget can be allocated across a model’s reasoning structure. We model generated claims first as a tree and then as a directed acyclic graph, with each claim assigned a hidden ground-truth label and a self-reported confidence score. We compare random, confidence-based, dependency-aware, dynamic-programming, and submodular Monte Carlo greedy verification strategies, along with adaptive threshold policies that estimate a claim’s potential downstream damage. Testing against nine real LLMs shows that models are systematically overconfident on false claims, while true/false confidence gaps are narrow and sometimes indistinguishable. These results suggest that structure-aware verification can outperform confidence-based checking and offers a practical alternative for verifying complex LLM reasoning.

Neural Mass Dynamics: From Critical Phase Transitions to Reaction-Diffusion Parameter Modulation in Alzheimer’s Tau Propagation

Arda Alisan, Rafia Hossain, Sophie Zhang

Currently, no mathematical account fully explains why Alzheimer’s tau pathology first targets regions such as the transentorhinal and entorhinal cortex. Hub degree, the leading structural explanation, is confounded with dynamical centrality, so connectivity alone cannot distinguish whether pathology follows a region’s wiring or its dynamics. We rank brain regions by their dynamical resistance to destabilization by linearizing a Wilson–Cowan neural mass model near criticality on an 83-region HCP connectome and using network control theory to compute each region’s minimum control energy. Across 685 ADNI subjects, tau accumulated preferentially in these hardest-to-tip regions, showing the opposite pattern from amyloid, and the relationship strengthened after controlling for hub degree. We then use this ranking in a personalized Fisher–Kolmogorov reaction-diffusion model calibrated to 43 subjects’ longitudinal tau-PET trajectories. Our simulations suggest that suppressing local tau growth outperforms suppressing diffusion, and that protecting regions with the highest baseline tau burden produces greater reductions in simulated tau spread than targeting regions by structural or dynamical vulnerability. These results suggest that burden-informed targeting may better guide future anti-tau interventions.

Private Searches, Public Patterns: Discovering Shapes of Attention with Google Trends

Xinrui Chen, Meriem Fereva, Aine Reeves

Online search behavior provides a rich signal of collective attention, and Google Trends is widely used to study these patterns. Because Google Trends requires search terms to be specified in advance, most studies rely on predefined, topic-specific dictionaries. We instead build a word bank from Google’s Top and Rising Queries and analyze the resulting time series using Symbolic Aggregate approximation (SAX). We then group search terms by the shape of their temporal trajectories using hierarchical clustering. The resulting clusters reveal interpretable categories of attention dynamics, including seasonal interest, gradual growth, and sudden bursts of activity. These recurring attention signatures can be identified across events and time periods and may support downstream applications such as market monitoring and forecasting.

Dynamic Fair Value Estimations of Prediction Market Contracts Using Financial Data

Manni Lin, Solon Sun, Lara Urteaga

This project investigates whether traditional financial market data can be used to estimate fair value in prediction markets. While existing work often uses prediction markets to inform financial market strategies, we examine the inverse relationship: whether financial time series can provide useful signals for prediction-market pricing. We develop allocation strategies along two pathways: single-series allocation and multi-series allocation. The single-series approach uses Sharpe ratio optimization to guide allocation, while the multi-series approach uses a dynamic portfolio informed by Granger causality and fair value estimates derived from financial signals. Together, these methods provide a framework for studying information flow between traditional financial markets and prediction markets and for identifying when prediction-market contracts may deviate from estimated fair value.

ComfiMap: Measuring Comfortable Walkability Through Pedestrian Thermal Comfort in New York City

Cathy Deng, Sophie He, Sakhi Kulkarni

Walking is one of the most common ways people move through cities, but rising temperatures and the urban heat island effect are making pedestrian travel increasingly uncomfortable and, at times, unsafe. Existing navigation tools typically optimize for distance, while many heat-aware approaches focus only on shade or surface air temperature. We propose ComfiMap, a framework for measuring comfortable walkability by combining pedestrian travel time with thermal comfort conditions along city streets. Using Lower Manhattan as a case study, we integrate microclimate models for mean radiant temperature, wind, and the Universal Thermal Climate Index (UTCI) to construct a high-resolution pedestrian network weighted by both distance and cumulative thermal discomfort. We then compare possible walking routes and summarize the trade-off between shorter and more comfortable paths over time. This framework could provide a basis for defining a comfort-walkability measure for an urban area, capturing how easily pedestrians can avoid excessive heat exposure while walking. Because the model relies on increasingly available urban surface data, ComfiMap can be extended to other cities and used to support heat-aware routing, urban design, and climate-adaptive planning.