



**Program for Research in
Mathematics, Engineering and Science**



Photo: Slava Gerovitch

**Sixteenth Annual Fall-Term
PRIMES Conference
October 24-25, 2026**

Saturday, October 24

Mathematics

Room 4-370, MIT

9:00-9:15 am Welcoming Remarks

- Prof. Michel Goemans, Head of the MIT Mathematics Department
- Prof. Pavel Etingof, PRIMES Chief Research Advisor
- Dr. Slava Gerovitch, PRIMES Program Director

9:15-10:45 am Session 1

- Shengdao (Bill) Ke and Zifang (Eric) Wu, "Effective Transcription for Low Affinity Binding sites" (mentors David Darrow, MIT, and Joseph Landsittel, Northwestern)
- Katrina Liu, "Higher-Order Co-Infection on Hypergraphs" (mentor Prof. Laura Schaposnik, University of Illinois Chicago)
- Edward Zhang, "Optical Emulation of Higher-Order Correlation Functions" (mentor Dr. Georgi Gary Rozenman)
- Eric Min, Arjun Pandey, and Ziyi (Kaya) Yang, "Implementing Fast Khovanov Homology in C With Fixed-Strand Combinatorial Reductions" (mentor Dr. Ryan Maguire, Draper)
- Dell Li, Maiya Qiu, Yutong Zhao, "New Methods towards High-Resolution Reconstruction of Saturn's Rings" (mentor Dr. Ryan Maguire, Draper)

11:15 am -12:30 pm Session 2

- Michael Luo, "A Tour of the Generalized Catalan Numbers" (mentor Dr. Xinchun Ma)
- Neil Kolekar, "On Support of Temperley-Lieb Immanants of Ribbon Decomposition Matrices" (mentor Son Nguyen)
- Ryan Bansal, "A Forest Analogue of Bruhat Order" (mentor Dora Woodruff)
- Rohith Thomas, "Injectivity of Symmetric Polynomial Maps on Partitions" (mentor Katherine Tung, Harvard)

- Michael Zhao, "Exact Promotion Order of All Two-Row Standard Young Tableaux" (mentor Laura Pierson, Harvard)

1:40-3:00 pm Session 3

- Tejo Madhavarapu, "Fair Duels after a False Start" (mentor Dr. Tanya Khovanova)
- Zack Berchenko "On the Cardinality of the Smith Set with Many Alternatives" (mentor Kaizhao Liu)
- Anderson Hao, "Limit Laws for Coefficients of Polynomial Powers" (mentor Dr. Dmitry Krachun, Princeton)
- Andrew Chai and Richard Wang, "Mixing Times of Toggle Markov Chains" (mentor Rupert Li, Stanford)
- Dou Lai, "An All-Sky Search for Moving Gamma-Ray Sources with Fermi-LAT" (mentor Prof. James Unwin, University of Illinois Chicago)

3:20-4:20 pm Session 4

- Grisham Paimagam, "The Union of Two Hamiltonian Paths or Cycles" (mentor Dr. Rikhav Shah)
- Minkyu Jo, "Ricci Flow on ALH Manifolds" (mentor Isaac M. Lopez, University of Chicago)
- Oren Akresh, "Weak Jordan Algebras from the Verlinde Category in Characteristic 5" (mentor Andrew Riesen)
- Connor Leong, "A Combinatorial Bijection Between Orbits on Flags" (mentor Kenta Suzuki, Princeton)

4:40-6:00 pm Session 5

- Manit Gupta, Terence Wu, and Alex Zhan, "A Topological Interpretation of Divisibility" (mentors Prof. Jim Coykendall, Clemson University, and Prof. Jared Kettinger, Grove City College)
- Mingyang Cen, Eric Guo, and Patrek Ragnarsson, "On the Ascent of Divisibility Properties to Monoid Algebras" (mentor Dr. Felix Gotti)
- Keshav Karumbunathan, Hyun-Jin Kim, and Vincent Wang, "Multiplicative Structure of Rational Monogenic Semidomains" (mentor Dr. Felix Gotti)

- Ishaan Mittal, Luca Pieleanu, and Shunyao Yan, "Atomic Density and Local Behavior of Irreducible Polynomials" (mentor Dr. Harold Polo, Clemson University)

Bioinformatics (in parallel with Math)

Room 4-270, MIT

2:00-3:00 pm Session 6

- Samuel Zhang, "SilentMethyl: A Multimodal Framework for CpG Methylation Prediction and Local Variant-Effect Inference" (mentors Prof. Gil Alterovitz, Harvard Medical School, and Dr. Shaojun Pei, Brigham and Women's hospital)
- Jerry Xu, "A Unified Protein Embedding Model with Local and Global Structural Sensitivity" (mentors Prof. Gil Alterovitz, Harvard Medical School, and Dr. Shaojun Pei, Brigham and Women's hospital)
- Lev Smolsky, "Localized Breath and Cough Acoustics Complement Participant Metadata in a Synchronized Tuberculosis Cohort" (mentor Prof. Alterovitz, Harvard Medical School)
- Rajarshi Mandal, "Calibrated Benchmarking and Adaptive Learning for ITN Allocation in Malaria Control" (mentor Prof. Gil Alterovitz, Harvard Medical School)

3:30-4:20 pm Session 7

- Michael Jian, "SpaCoEx: Sparse Gene Selection for Spatially Varying Co-expression in Spatial Transcriptomics" (mentor Prof. Gil Alterovitz, Harvard Medical School)
- Aadya Goel, "Multiscale Transcriptomic Architecture of the ALS Motor Cortex" (mentor Prof. Gil Alterovitz, Harvard Medical School)
- Xuanye Zhang and Anirudh Pulugurtha, "Computational Characterization of Mean Arterial Pressure Trajectories After Spinal Anesthesia for Cesarean Delivery" (mentors Prof. Vesela Kovacheva and Dr. Ricardo Kleinlein, Harvard Medical School and Brigham and Women's Hospital)

Sunday, October 25

Mathematics

Room 4-370, MIT

9:00-10:00 am Session 8

- Timothy Chen and Max Levin, “Autoregressive ODE Discovery from Data” (mentors Prof. Lu Lu, Yale, and Prof. Yixiang Deng, University of Delaware)
- Alex Liu and Anthony Zou, "Efficient Catalytic Algorithms for All-Pairs Shortest Paths" (mentor Zoe Xi)
- Omar Graia, Caroline Huang, and Anna Zhou, “Probabilistic Fault-Tolerant Metric Dimension” (mentor Prof. Jesse Geneson, San José State University)

10:30-11:30 am Session 9

- Jiaheng (Gary) Li, "On Signature Characters and Unitary Loci for Type A_1 , A_2 Quantum Groups at Roots of Unity" (mentor Frank Wang)
- Harry Gao, "Semisimplification of the Coordinate Ring of $GL(m)_{J_1} \oplus nJ_{p-1}$ in Positive Characteristic" (mentor Dr. Arun Kannan, DRW)
- Derek Aoki, "On Open Borel Orbits in Conormal Bundles to Schubert Cells" (mentor Hamilton Wan)
- Jessica Yan, "Computing Isomorphisms Between Groups Defined by Integer Matrices" (mentor Prof. Maria Sabitova, CCNY – CUNY)

12:00-12:50 pm Session 10

- Burak Aricioglu, "A Closed Formula for Eigenweights" (mentor Zeyu Wang)
- Wesley Lin, "A Representation-Theoretic Computation of the Alexander Polynomial of Satellite Knots" (mentor Kenta Suzuki, Princeton University)
- Leo Hong and Ari Shpilman, "Orbital Integrals and Shalika Germs" (mentor Dr. Thomas Rüd)

2:00-3:05 pm Session 11

- Brian Sun and Seabert Mao, "Smoothness Characterizations and Duality of the Type II Theta Correspondence and Basic Fourier–Jacobi Models" (mentor Dr. Jialiang Zou)
- Grant Blitz, "Coprimality and Divisibility in Integer Matrices" (mentor Prof. Maxim Arnold, University of Texas Dallas)
- Niranjan Rao, "Cubic Surfaces over F_q with 27 Lines over F_{q^2} " (mentor Xinyu Fang, Harvard University)
- Abhigyan Singh, " T -equivariant Modules over the Ring $k[x_1, x_2, \dots, x_n][y_1, y_2, \dots, y_{n-1}]/(x_i y_i)_{i=1}^{n-1}$ " (mentor Dr. Linus Hamann, Harvard)

3:25-4:10 pm Session 12

- Madeline Hu, "The Associative Zesting of an Exact Factorization of Fusion Categories" (mentors Prof. Julia Plavnik and Dr. Héctor Peña Pollastri, Indiana University Bloomington)
- Tarun Rapaka, "Unfactorizable Zestings of Deligne Products of Tambara–Yamagami Fusion Categories" (mentors Prof. Julia Plavnik and Dr. Héctor Peña Pollastri, Indiana University Bloomington)
- Dannan Li, "A Bound on the Fractal Uncertainty Principle for a Square Alphabet" (mentor Alain Kangabire)

4:30-5:30 pm Session 13

- Charlie Gu and Kritika Chandra, "Relative Humidity and Wind Speed for Long-Term Probabilistic Load Forecasting" (mentor Dr. Tao Hong, University of North Carolina at Charlotte)
- Aarush Kulkarni, Sirui (Kris) Sheng, and James Wu, "Directional Curvature for Efficient Word2Vec Optimization" (mentor Dr. Willy Rodriguez, Torus AI)
- Yunong (Melody) Wu, Gary Shen, and Tanish Parida, "Predicting Cognitive Decline from Baseline Assessments: An OASIS-3 Study" (mentor Dr. Marly Gotti, Arcus Biosciences)

2026 PRIMES FALL-TERM CONFERENCE ABSTRACTS

SATURDAY, OCTOBER 24

SESSION 1

Shengdao (Bill) Ke and Zifang (Eric) Wu

Effective Transcription for Low Affinity Binding sites

Mentors: David Darrow, MIT, and Joseph Landsittel, Northwestern

Multicellular eukaryotes often regulate transcription through clusters of low-affinity binding sites rather than single high-affinity sites. Such architectures may help cognate transcription factors distinguish functional regulatory regions from the many decoy sequences distributed throughout large genomes. However, clustering weak sites introduces a competing challenge: non-cognate transcription factors can also bind to the same regulatory region, reducing specificity through competition with the cognate factor. Here, we study these binding-site clusters using an information-theoretic framework, focusing on how binding-site number and affinity affect the reliability of information transmission from transcription-factor concentration to gene expression. Numerical optimization shows that information transmission depends jointly on site number and binding affinity. We find that at higher transcription-factor concentrations, promoters with multiple binding sites outperform single-site architectures, providing a possible information-theoretic explanation for the prevalence of clustered regulatory regions in eukaryotic genomes.

Katrina Liu

Higher-Order Co-Infection on Hypergraphs

Mentor: Prof. Laura Schaposnik, University of Illinois Chicago

People encounter infections through overlapping groups, such as households, classrooms, and workplaces. Understanding how these shared contacts affect the spread of multiple pathogens is critical to determining what information epidemic models need to retain. We introduce HyCo-SIS (Higher-order Co-infection SIS), a discrete-time SIS model on hypergraphs with explicit reinforcement between overlapping active groups and pairwise transmission. We investigate how this reinforcement affects infection dynamics and whether simplified representations of hypergraphs preserve the resulting behavior. We derive an exact representation of overlap reinforcement as target-dependent four-person interactions and determine when a symmetric representation is possible. Numerical experiments examine reinforcement strength, contact structure, and susceptibility to secondary infection, as well as calibrated pairwise approximations. Our results demonstrate that redistributing overlap weights can substantially change finite-time co-infection even when every individual's total overlap coefficient is preserved. We further derive an exact five-person calculation and demonstrate why equal mean infection pressure does not necessarily produce equal acquisition probabilities. These findings identify how the allocation of higher-order interactions can influence contagion dynamics beyond their aggregate strength. Beyond epidemiology, this work contributes to the statistical physics of spreading processes by examining how the arrangement of many-body interactions shapes collective dynamics, and what information is lost when those interactions are simplified.

Edward Zhang

Optical Emulation of Higher-Order Correlation Functions

Mentor: Dr. Georgi Gary Rozenman

Optical correlation measurements provide a powerful framework for exploring photon statistics and quantum coherence, with higher-order correlations carrying information about a light source that no measurement of first- or second-order coherence can supply. Here, we present a classical optical emulation apparatus of the fourth-order correlation function $g^{(4)}$, built with a pulsed coherent laser, programmable polarization routing, and camera detectors to reproduce selected coincidence-suppression signatures of single-photon emission. We prove that, under ideal binary routing, the normalized four-fold coincidence on a fixed-anchor delay slice is unity except for five lines on which two detection times coincide, where it vanishes—the same coincidence-exclusion topology that a single-photon emitter imposes—and that this result holds for any pulse-energy statistic. Experimentally, we observe a suppressed zero-delay second-order correlation, three coincidence-suppression lines in third-order delay maps, and all five predicted lines in fixed-anchor fourth-order maps. Empirical models constructed from measured pair correlations reproduce the third- and fourth-order maps with Pearson correlations of 0.997 and 0.987. Because the exclusion structure originates in the routing rather than in the photon statistics, the platform separates the geometry of higher-order coincidence measurements from the quantum nature of the source and offers an inexpensive test bed for detector, noise and estimator effects in multi-detector correlation experiments. The platform provides a practical setting for studying higher-order correlations, noise, and correlation estimators, with applications in quantum optics education and experimental method development.

Eric Min, Arjun Pandey, and Ziyi (Kaya) Yang

Implementing Fast Khovanov Homology in C With Fixed-Strand Combinatorial Reductions

Mentor: Dr. Ryan Maguire, Draper

Khovanov homology is a knot invariant that categorifies the Jones polynomial, but direct computation is limited by the 2^n states in the cube of resolutions. Bar-Natan's Fast Khovanov homology algorithm reduces this cost in practice by constructing the complex locally and repeatedly simplifying it using delooping and Gaussian elimination. We implement this scanning approach in C, based in part on the architecture of JavaKh but using concrete C representations for the smoothings, cobordisms, linear combinations, and differential matrices required by the computation. The program constructs and simplifies tangle complexes one crossing at a time, supports integral and finite-field coefficients, and accepts braid, planar-diagram, and Dowker–Thistlethwaite inputs through a common pipeline. We validate the implementation against independently tabulated Khovanov homology. The C implementation is intended to reduce memory and runtime overhead and to support computations on a larger scale, and we evaluate its performance relative to existing implementations. One target application is to extend computations to prime knots with 18–20 crossings and, where feasible, beyond, while comparing the distinguishing power of Khovanov homology with that of the Jones and HOMFLY-PT polynomials.

Dell Li, Maiya Qiu, and Yutong Zhao

New Methods towards High-Resolution Reconstruction of Saturn's Rings

Mentor: Dr. Ryan Maguire, Draper

This talk focuses on the reconstruction of Saturn's ring profiles from radio occultation data. We find that increasing the spatial resolution from 1 km (the resolution of existing reconstructions) to 50 m or 20 m reveals additional roots whose contributions are not accounted for in the stationary phase approximation (SPA) step of the reconstruction. We therefore discuss several numerical root-finding and continuation methods for tracking all relevant solution branches. We also provide an overall framework for the SPA step by introducing a branch-bookkeeping layer to track root branches near bifurcations. Finally, we propose a residual-validation benchmark that compares the ordinary stationary point sum with a full angular integral.

Michael Luo

A Tour of the Generalized Catalan Numbers

Mentor: Dr. Xinchun Ma

We define in sequence the Catalan numbers, the q -Catalan numbers, the q, t -Catalan numbers, the rational q, t -Catalan numbers, and finally the generalized q, t -Catalan numbers. We characterize a superset of the support of the generalized q, t -Catalan numbers in terms of a convex hull. We also show the existence of stable diagonals with equal coefficients of constant degree, generalizing a theorem of Lee–Li.

Neil Kolekar

On Support of Temperley–Lieb Immanants of Ribbon Decomposition Matrices

Mentor: Son Nguyen

A 2005 result of Rhoades and Skandera establishes that Temperley–Lieb immanants, a subset of the dual canonical basis, are Schur-positive when evaluated on Jacobi–Trudi matrices. In 2025, Nguyen and Pylyavskyy established a Littlewood–Richardson rule for Temperley–Lieb immanants of Jacobi–Trudi matrices via *shuffle tableaux*. This rule was later extended to Temperley–Lieb immanants evaluated on *ribbon decomposition matrices*, a class of matrices which subsume several classical determinantal formulas for skew Schur functions, including the Jacobi–Trudi, Giambelli, and Lascoux–Pragacz identities.

In 2026, Yevtushenko proved that for a restricted class of shapes, the Schur supports of Temperley–Lieb immanants of the corresponding Jacobi–Trudi matrices have a maximal element with respect to the dominance order. He conjectured that the Schur support contains both maximal and minimal elements with respect to the dominance order for all shapes. In this talk, we resolve and generalize Yevtushenko’s conjecture by proving that the Schur supports of Temperley–Lieb immanants of all ribbon decomposition matrices contain maximal and minimal elements, each appearing with coefficient one.

Ryan Bansal

A Forest Analogue of Bruhat Order

Mentor: Dora Woodruff

In Schubert calculus, polynomial multiplication encodes geometric intersection problems. Monk’s rule, which describes multiplication by a linear Schubert polynomial, traces out the Bruhat order, a partial order on permutations. Forest polynomials, introduced by Nadeau and Tewari in 2023, are indexed by binary forests rather than permutations and satisfy an analogous multiplication rule.

We use this rule to define the forest Monk poset and ask which properties of Bruhat order persist. Through examples, I will present two answers. First, we classify and enumerate the principal lower intervals of this poset that are Boolean lattices. These are exactly the intervals below forests that avoid bent corners in their internal paths. This parallels Tenner’s pattern-avoidance characterization for permutations. Second, an explicit chain decomposition proves that every finite window of the poset is strongly Sperner: a union of r families of mutually incomparable forests is no larger than the r largest ranks combined. A classical proof of the corresponding property for Bruhat order uses the hard Lefschetz theorem; our proof is entirely combinatorial. I will conclude with a shellability conjecture and the search for geometric explanations of these results.

Rohith Thomas

Injectivity of Symmetric Polynomial Maps on Partitions

Mentor: Katherine Tung, Harvard

Introduced by Ballantine, Beck, and Merca, the elementary symmetric partition function pre_k , defined on the set of partitions with at least k parts, has been a topic of recent interest. We prove that pre_k is injective on the set of m -ary partitions for positive integers $m \geq k$, generalizing a $k = 2$ result of Ballantine, Beck, Merca, and Sagan, and complementing a result of Hadelyn, Niergarth, [Weiyu] Li, and [Wenhui] Li showing that, for each $k \geq 3$, pre_k is not injective on partitions of n with length $2k$ for infinitely many n . We introduce the skew Schur partition function $\text{prs}_{\lambda'/\mu'}$, and prove injectivity results for particular choices of λ', μ' .

Michael Zhao

Exact Promotion Order of All Two-Row Standard Young Tableaux

Mentor: Laura Pierson, Harvard

Promotion is a bijective operation on standard Young tableaux of any fixed shape. The orbit length, the number of times promotion needs to be applied to a tableau to return to the original state, is well-studied in certain cases, such as rectangular tableaux. However, the orbit length of promotion for other Young tableaux is less well-understood. We determine this orbit length for Young tableaux with two rows. To do so, we explore structures constructed from Young tableaux which reveal information about the promotion operation, and track the state of the tableaux using these structures to characterize the promotion order.

SESSION 3

Tejo Madhavarapu

Fair Duels after a False Start

Mentor: Dr. Tanya Khovanova

How do we make a duel fair? One way to do this is to only allow the player who's currently at a disadvantage to shoot next. When we do this for two lousy shooters who are equally skilled, we get a shooting sequence that doesn't repeat (and is therefore annoying to work with in an actual duel). But what happens if there's a mistake in determining who shoots next? Surprisingly, this gives us a sequence that eventually starts repeating! In this talk, we explore exactly why that happens.

Zack Berchenko

On the Cardinality of the Smith Set with Many Alternatives

Mentor: Kaizhao Liu

For some integer $\ell \geq 2$, consider $m = 2\ell - 1$ voters, each of whom has a preference ranking over n alternatives. The Smith set is the smallest nonempty set of alternatives that defeat every alternative outside the set in pairwise majority comparisons. In general, the size of the Smith set serves as a structural measure of consensus: a small set identifies a clear top tier of alternatives, while a large set indicates widespread cyclic preferences. Traditionally, voting theory studies a regime in which a large, growing population of voters ($\ell \rightarrow \infty$) evaluates a small, fixed number of alternatives (n). However, recent applications necessitate studying a different regime in which ℓ is fixed and $n \rightarrow \infty$. For example, in the alignment of large language models via Reinforcement Learning from Human Feedback, a fixed panel of human expert labelers (voters) provides pairwise comparisons over a massive pool of potential model

outputs (alternatives). Under an impartial culture and the Luce model, two common probability distributions for ranked voting, we characterize the asymptotic distribution of the Smith set's cardinality for fixed ℓ as $n \rightarrow \infty$.

Anderson Hao

Limit Laws for Coefficients of Polynomial Powers

Mentor: Dr. Dmitry Krachun, Princeton

We study the distribution of coefficients in the expansion of $P(x)^n$ as $n \rightarrow \infty$ for real polynomials $P(x)$. Graphing the coefficient distribution reveals various peaks of eccentric shapes and of different widths. Some peaks oscillate, with a broad range in oscillation frequencies, while other peaks are not even symmetric. Our primary goal is to understand the visual phenomena in these graphs.

When $P(x)$ has only nonnegative coefficients, the limiting distribution is a single Gaussian peak, which can be seen as a variant of the Central Limit Theorem. For the general case, richer structure appears, and by carefully analyzing certain complex integrals over the unit circle, we identify possible limiting shapes of the peaks with a countable collection of universality classes indexed by integers $q \geq 2$.

For $q = 2$, we recover the Gaussian peak. Interestingly, for $q = 3$ the limiting shape shares some features with the Tracy-Widom law famously appearing in the random matrix theory. For each $q \geq 2$ we prove that the local peak converges to the corresponding limiting distribution when renormalized by $n^{1/q}$, generalizing the fact that fluctuations in the Central Limit Theorem are of order $n^{1/2}$.

In addition, we explain the emergence of oscillations and develop formulae for predicting oscillation frequency in certain cases. We also establish a bound on the number of peaks of each order, present the results of computer experiments, and pose a conjecture on the possible peak combinations.

Andrew Chai and Richard Wang

Mixing Times of Toggle Markov Chains

Mentor: Rupert Li, Stanford

The notion of *rowmotion*, a bijective operator on the distributive lattice $J(P)$ of order ideals of a finite poset P , was first defined in 2012 by Striker and Williams. Since then, a substantial body of work has explored both the algebraic and combinatorial properties of rowmotion. In 2023, Defant, Li, and Nestoridi defined *rowmotion Markov chains* $\mathbf{M}_{J(P)}$ by considering stochastic transitions on $J(P)$. Building on Striker's generalization of rowmotion to the toggle operator, they further introduced *toggle Markov chains* and initiated the study of mixing times in such chains. One of their results concerning the mixing of toggle Markov chains on Boolean lattices is generalized in this work.

In this talk, our main results address an open question posed by Defant, Li, and Nestoridi, concerning the mixing of two natural classes of toggle Markov chains $\mathbf{T}$ which we term *lower- k* and *upper- k* chains, where the integer k pertains to a constraint imposed on the state space of $\mathbf{T}$. Respectively conditioned on k or $n - k$ being constant, where $n = |P|$, we show the mixing time is $\Theta(n^2)$ for both chains. For more general k , we provide broader asymptotic bounds.

In our work, we also discuss an interesting connection between bounded- k toggle chains and a variant of the stochastic box-ball system. The box-ball system is a discrete cellular automaton introduced by Takahashi and Satsuma in 1990, and it was first considered in a stochastic setting by Keating, Kim, Loeser, and Lyu in 2025. The box-ball interpretation of bounded- k chains provides an intuitive framework under which parts of our main results are proven.

Dou Lai

An All-Sky Search for Moving Gamma-Ray Sources with Fermi-LAT

Mentor: Prof. James Unwin, University of Illinois Chicago

Nearby compact gamma-ray sources can move appreciably during the Fermi-LAT mission and thus evade searches that assume a stationary sky. Compact dark matter halos, dark matter spikes around primordial black holes (PBH), and evaporating PBH provide potential examples. We perform a novel all-sky search for steady point sources moving at $0.25^\circ/\text{yr}$ – $0.8^\circ/\text{yr}$ in ten years of 1-100 GeV Fermi-LAT data. For a transverse speed of 200 km s^{-1} , this angular range probes distances of roughly $3.0 \times 10^3 \text{ AU}$ to $9.7 \times 10^3 \text{ AU}$. We find 577 candidate tracks, consistent with 533 ± 27.0 tracks in time-scrambled skies. We place a model-independent 90% upper limit of 80.8 recovered signal tracks. At peak calibrated sensitivity, this corresponds to fewer than 123 steady, full-sky sources with an incident E^{-2} photon spectrum. We give the calibrated detection efficiency of the search to steady point sources and use it to derive illustrative constraints on both evaporating PBHs and PBHs surrounded by annihilating dark matter halos.

SESSION 4

Grisham Paimagam

The Union of Two Hamiltonian Paths or Cycles

Mentor: Dr. Rikhav Shah

Each permutation σ on $[n]$ corresponds naturally to a Hamiltonian path: connect $\sigma(j)$ to $\sigma(j+1)$ for each $j \in [n-1]$. To obtain a Hamiltonian cycle, connect $\sigma(n)$ back to $\sigma(1)$. This talk considers graphs which result from taking the union of the Hamiltonian path or cycle for the identity permutation, and for a handful of permutations arising in number theory. We find manifold-like structure in these graphs, and formally prove topological and metric embeddings into particular manifolds.

Minkyu Jo

Ricci Flow on ALH Manifolds

Mentor: Isaac M. Lopez, University of Chicago

An asymptotically locally hyperbolic (ALH) manifold resembles a quotient of hyperbolic space outside a compact set. We show that, under suitable decay and regularity assumptions and up to rescaling, a Ricci flow started from an ALH metric preserves the ALH condition on finite time intervals where curvature remains bounded, making the notion of an ALH Ricci flow well-defined. We then study the evolution of geometric mass, an invariant determined by the geometry at infinity. Under the mass evolution law for normalized Ricci flow, initially nonzero mass decays exponentially and tends to zero precisely when the flow exists for all forward time. A nonzero limiting mass instead characterizes the formation of a finite-time singularity, linking the asymptotic geometry to the lifetime of the flow.

Oren Akresh

Weak Jordan Algebras from the Verlinde Category in Characteristic 5

Mentor: Andrew Riesen

The classical Kantor-Koecher-Tits (KKT) construction establishes a connection between Lie algebras and Jordan algebras. In characteristic 5, this connection takes an unexpected form—applying the construction to Lie algebras in the Verlinde category yields weak Jordan algebras that satisfy a weakened version of the Jordan identity. This talk combines that characteristic 5 phenomenon with the Freudenthal-Tits magic square and work of Elduque and Okubo to construct simple weak Jordan algebras associated with the exceptional Lie algebras $\mathfrak{f}_4$, $\mathfrak{e}_6$, $\mathfrak{e}_7$, and $\mathfrak{e}_8$. These algebras fail the full Jordan identity. An explicit 13-dimensional example built from Hermitian 3×3 matrices over k further illustrates these results, with its simplicity and failure of the full Jordan identity verified directly.

Connor Leong

A Combinatorial Bijection Between Orbits on Flags

Mentor: Kenta Suzuki, Princeton

How can two lines sit inside a plane? They can be parallel, intersect, or coincide. Questions like this, about the ways subspaces can be arranged, are naturally phrased in terms of group actions on vector spaces and the orbits they form. In this talk, we study the arrangement of complete flags, nested chains of subspaces each one dimension larger than the last. A conjecture arising from the relative Langlands program predicts that two very different groups partition flags into sets of orbits in bijection with each other. By encoding each orbit in a combinatorial diagram, we confirm this prediction for one dual pair and give an explicit bijection between the two sets.

SESSION 5

Manit Gupta, Terence Wu, and Alex Zhan

A Topological Interpretation of Divisibility

Mentors: Prof. Jim Coykendall, Clemson University, and Prof. Jared Kettinger, Grove City College

We study the divisibility structure of a commutative ring through the strongly multiplicatively closed (SMC) topology defined over the ring, whose open sets are precisely the subsets closed under multiplication by arbitrary ring elements. This Alexandrov topology has a basis of principal ideals and extends naturally to commutative semigroups. We show that compact open sets are precisely finite unions of principal ideals and that a unital commutative ring is Noetherian as an SMC space if and only if it is a Cohen–Kaplansky ring. We also obtain criteria for when the entire ring is compact under the topology and when the topology is Artinian. The characterization of compactness for open sets leads to a study of finite unions of ideals and principal ideals.

Mingyang Cen, Eric Guo, and Patrek Ragnarsson

On the Ascent of Divisibility Properties to Monoid Algebras

Mentor: Dr. Felix Gotti

Let $R[M]$ denote the monoid algebra of a cancellative, commutative, and torsion-free monoid M over an integral domain R . In this talk, we study the ascent of certain divisibility properties from a monoid to its algebra. We first investigate the relationship between the pre-Schreier and maximal common divisor properties of monoids, as well as the ascent of finiteness properties to monoid algebras. We also say that an integral domain satisfies the D-property if, for all nonzero elements a, b, c , the conditions that a divides bc and that a and b have no nonunit common divisor imply that a divides c . Finally, we explicitly characterize the ascent of the D-property to monoid algebras over integral domains R and Puiseux monoids M , stating exactly when $R[M]$ satisfies the D-property.

Keshav Karumbunathan, Hyun-Jin Kim, and Vincent Wang

Multiplicative Structure of Rational Monogenic Semidomains

Mentor: Dr. Felix Gotti

For $q \in \mathbb{Q}_{>0}$, let $S_q := \mathbb{N}_0[q]$. We investigate divisibility and factorization in the multiplicative monoid S_q^* . We prove that S_q^* is Furstenberg if and only if its technical monoid T_q is Furstenberg. We construct explicit families of non-Furstenberg and non-atomic rational monogenic semidomains. For $q < 1/2$ outside the ACCP cases, we characterize atoms and ACCP elements through unique polynomial representations and prove that the MCD and quasi-ACCP properties are equivalent. We also obtain MCD-finiteness criteria and construct pairs with arbitrarily many nonassociate MCDs in every nonfactorial case. When every prime-power factor of $d(q)$ exceeds $n(q)$, the monoid S_q^* is not Furstenberg, and it is MCD-finite but neither MCD nor quasi-ACCP. Finally, in the non-ACCP range, we characterize ACCP elements by a uniform bound on the coefficients of their polynomial representations.

Ishaan Mittal, Luca Pieleanu, and Shun Yao Yan

Atomic Density and Local Behavior of Irreducible Polynomials

Mentor: Dr. Harold Polo, Clemson University

We study the asymptotic and constructive behavior of irreducible polynomials in the semiring $\mathbb{N}_0[x]$ of polynomials with nonnegative integer coefficients. First, we strengthen the iterated-limit formulation of atomic density by proving that the proportion of irreducibles in uniform-height boxes tends to 1 as both the degree and coefficient-height bounds tend to infinity. For fixed degree $d \geq 1$, we extend this analysis to nonuniform coefficient boxes and show that the asymptotic proportion of irreducibles is $1/\zeta(d+1)$ as every coefficient bound tends to infinity. Building on these density results, we study Goldbach-type decompositions of polynomials into sums of two multiplicative atoms, obtaining an asymptotic formula for the number of ordered decompositions as all coefficients tend to infinity for fixed degree. Finally, we investigate Turán-type bounded-distance problems, studying how much a polynomial's coefficients must be changed in coefficientwise ℓ_1 distance to obtain an irreducible polynomial of the same degree. We establish that all monic cubics with positive constant term are within distance 1 of an irreducible, we prove a general degree-dependent upper bound of $d/2 + O(\sqrt{d \log d})$, and we show degree-independent bounds for polynomials with bounded interior coefficients or at most $\frac{1}{2} \log_2 d$ coefficients greater than one. In particular, for polynomials with at most $\frac{1}{2} \log_2 d$ coefficients greater than one, we obtain a distance bound of 9 for degrees $d \geq 4096$. We further investigate bounded-distance problems for more general classes of polynomials.

SESSION 6: BIOINFORMATICS

Samuel Zhang

SilentMethyl: A Multimodal Framework for CpG Methylation Prediction and Local Variant-Effect Inference

Mentors: Prof. Gil Alterovitz, Harvard Medical School, and Dr. Shaojun Pei, Brigham and Women's hospital

Predicting how nucleotide variants influence local methylation remains challenging. We developed SilentMethyl, a multimodal framework integrating DNA sequence with epigenomic context to predict methylation and infer local variant effects. SilentMethyl achieved a beta-value MAE of 0.0914 and ROC-AUC of 0.9744 across 26,570 held-out CpGs, outperforming sequence-only and context-only models as well as retrained CpGenie and DeepCpG. Context reduced prediction error by 22.4% and 22.9% at CpGs with high ATAC-seq or H3K27ac signal, respectively. Cross-tissue predictions were most accurate at CpGs with low methylation variability across tissues. Predicted effects showed modest concordance with independent mQTL and allele-specific methylation data, with signed Spearman correlations of 0.236 in eGTEx breast and 0.149 in GENOA blood. Variant-CpG distance contributed substantially

to apparent discrimination of mQTL-associated variants, with fusion AUROC decreasing from 0.600 to 0.556 after distance matching. Context influenced variant predictions, but fusion did not clearly outperform sequence-only on external validation benchmarks. SilentMethyl also predicted a local methylation change around a pathogenic STK11 stop-gain variant and independently prioritized a synonymous NCOA2 variant among 440 held-out TCGA-BRCA SNV-CpG pairs. Overall, SilentMethyl provides a multimodal framework for accurate CpG methylation prediction and inference of local variant effects, enabling variant prioritization for targeted experimental validation.

Jerry Xu

A Unified Protein Embedding Model with Local and Global Structural Sensitivity

Mentors: Prof. Gil Alterovitz, Harvard Medical School, and Dr. Shaojun Pei, Brigham and Women's hospital

Protein structural comparison is key to many research tasks, including evolutionary analysis, functional annotation, and the interpretation of missense variants. With the rapid growth of known and predicted structures in protein databanks, we need more efficient and comprehensive comparison methods. Traditional superposition-based structure alignment tools, such as TM-Align or DALI, are accurate but computationally expensive and impractical at scale. Existing protein language models (PLMs), such as TM-Vec, improve computational efficiency but only capture global structural similarity, overlooking important motif-level structural details. In this talk, we propose a novel PLM enabling efficient embedding-based structural comparison while also capturing both global and local structural similarity. Our model was trained on both TM-score, a global similarity metric, as well as PIS (patch isometry score), a local similarity metric derived from the Gromov-Wasserstein distortion between the 32-residue C α neighbourhoods of the two residues, which requires neither a superposition nor residue correspondences within the patch. Training occurred in two stages. Firstly, a Siamese E(3)-invariant graph network serves as a teacher, learning to embed each residue's structure patch to predict the PIS, and pools these embeddings to predict TM-score. Secondly, a student utilizing only sequences, a transformer over ESM-2 embeddings, is distilled from the frozen teacher. It learns to reproduce the teacher's per-residue and pooled embeddings, while also fitting the PIS and the TM-score. We built the training data from 985,230 pairs of Swiss-Prot proteins (351,443 proteins; 12,339,034 labelled residue pairs), grouped by 30%-identity sequence clusters. On held-out test clusters, the teacher achieves $R^2 = 0.76, r = 0.87$ for PIS and $R^2 = 0.93, r = 0.97$ for TM, while the student achieves $R^2 = 0.54, r = 0.76$ for PIS and $R^2 = 0.93, r = 0.96$ for TM. Our model fulfills two key roles: first, it rapidly detects global structural differences. Second, it supports fine-grained structural assessments, improving sensitivity to subtle but functionally important structural changes.

Lev Smolsky

Localized Breath and Cough Acoustics Complement Participant Metadata in a Synchronized Tuberculosis Cohort

Mentor: Prof. Alterovitz, Harvard Medical School

Most tuberculosis (TB) audio studies examine cough or breath separately, leaving the value of synchronized respiratory modalities unclear. We combined participant metadata, prompted cough and localized breath from 200 participants (91 with current pulmonary TB; 109 healthy comparison participants). In the primary nested analysis, low-capacity fusion achieved an area under the receiver operating characteristic curve (AUROC) of 0.7457 versus 0.7097 for recalibrated metadata, a difference of 0.0360 (95% confidence interval, 0.0062–0.0658), and improved log loss and Brier score. Full-pipeline permutation supported the log-loss improvement ($P = 0.032$); evidence for the AUROC improvement was weaker ($P = 0.290$). A post hoc development candidate achieved AUROC 0.7809 within the same cohort. Broader breath pooling reduced performance. External datasets lacked the synchronized inputs needed to evaluate the complete model and revealed acquisition sensitivity in its cough branch. These results provide internal evidence that localized breath and cough contribute complementary information beyond participant metadata.

Rajarshi Mandal

Calibrated Benchmarking and Adaptive Learning for ITN Allocation in Malaria Control

Mentor: Prof. Gil Alterovitz, Harvard Medical School

Malaria control programs distribute a fixed supply of insecticide-treated nets (ITNs) across regions, typically in proportion to population. Prior reinforcement learning work on malaria used fictional environments and rarely compared against this population-proportional baseline. We release an open benchmark based on a compartmental transmission model calibrated to real prevalence and incidence data from up to 50 subnational regions. PPO and a simple greedy planner both beat the baseline by up to about 6 percentage points of simulated infection reduction. When transmission is known the problem is near-myopic, and greedy allocation nearly matches an oracle planner. When transmission must instead be learned from surveillance, optimizing once against a prior performs worse than the baseline. Because distribution campaigns can also improve surveillance where they operate, every allocation is both an intervention and an experiment. We formulate adaptive allocation as constrained Bayesian dual control and show that certainty equivalence can stop learning about regions it underfunds. Our information-directed rule ARMOR-IDA also rewards the information each allocation yields. It achieves the largest gain of any adaptive method and outperforms certainty equivalence, Thompson sampling, and UCB in mean and worst-case reduction. After being calibrated in 2015, the simulator predicts prevalence through 2022 across 474 regions with correlation 0.76.

SESSION 7: BIOINFORMATICS

Michael Jian

SpaCoEx: Sparse Gene Selection for Spatially Varying Co-expression in Spatial Transcriptomics

Mentor: Prof. Gil Alterovitz, Harvard Medical School

Spatial transcriptomics enables gene expression to be measured while preserving tissue location, but most existing analyses focus on spatial variation in individual genes or expression-defined domains. Here, we introduce SpaCoEx, a sparse spatial representation framework that integrates gene-expression levels with spatially varying gene-gene co-expression. SpaCoEx first performs structured gene selection then uses the selected genes to construct both expression-level features and local co-expression features. We applied SpaCoEx to two human spatial transcriptomics samples. In the cutaneous squamous cell carcinoma dataset, SpaCoEx selected 29 of 45 keratinocyte-related genes while preserving 96.74% of the spatial co-expression variation. In the breast cancer dataset, SpaCoEx identified spatially varying co-expression between B2M and HLA-C, a biologically meaningful major histocompatibility complex (MHC) class I antigen-presentation gene pair. Their local correlation was significantly higher in cancer-associated regions than in non-cancer regions (mean difference = 0.30, spatially adjusted SE = 0.045, $P < 1.0 \times 10^{-10}$), whereas B2M and HLA-C expression individually did not differ significantly between cancer and non-cancer. In benchmarking against manual tissue annotations, co-expression-only SpaCoEx achieved the strongest spatial coherence (percentage of abnormal spots [PAS] = 0.076), while the joint expression/co-expression representation achieved the highest annotation agreement, with an adjusted Rand Index (ARI) of 0.584 and normalized mutual information (NMI) of 0.663. By integrating marginal gene-expression information with local gene-gene co-expression structure, SpaCoEx provides a sparse, low-dimensional, and interpretable representation of spatial transcriptomics data that captures complementary aspects of tissue organization beyond expression-based variation alone.

Aadya Goel

Multiscale Transcriptomic Architecture of the ALS Motor Cortex

Mentor: Prof. Gil Alterovitz, Harvard Medical School

Amyotrophic lateral sclerosis (ALS) exhibits marked clinical and molecular heterogeneity, but how this heterogeneity is organized within disease-relevant central nervous system regions remains incompletely understood. Because pooled cross-region analyses may confound regional variation with subtype structure, we focused on the motor cortex, a key site of ALS pathology, to resolve its internal transcriptomic organization. Using approximately 600 postmortem motor cortex transcriptomes from the Target ALS resource, we applied gene set variation analysis to generate sample-level pathway activity profiles and performed consensus clustering in pathway space. We identified seven reproducible ALS motor cortex subtypes characterized by distinct inflammatory, damage-response, metabolic, glial, and neuronal/synaptic programs. Donor-held-out projection recovered the seven subtype assignments with a mean balanced accuracy of 0.939, supporting the reproducibility of the pathway-defined subtype structure in unseen donors. Rather than forming isolated molecular classes, these subtypes occupied a structured pathway-space landscape with both discrete subtype-centered peaks and continuous biological gradients. Principal component analysis and representative pathway overlays defined a major axis from an upper inflammatory/damage region to a lower neuronal/metabolic/remodeling region, with C7 forming a neuronal/synaptic connector within the lower landscape. Co-expression network analysis further supported this multiscale organization, revealing family-level network backbones, subtype-specific modular refinements, and bridge-like rewiring. Metadata-informed analyses showed that this molecular architecture was not primarily organized by disease duration. Instead, the landscape was more clearly associated with clinical and genetic-background variables, including ALS/FTD family history and mutation-related signals. In particular, family-history-positive and C9orf72 repeat expansion-positive samples were more represented in the lower neuronal/metabolic/remodeling region, whereas family-history-negative samples and C9orf72-negative cases were more represented in the upper inflammatory/damage region. These findings support a model in which ALS motor cortex heterogeneity is organized as a multiscale transcriptomic architecture more closely associated with genetic-background variables than with disease duration. This framework provides a basis for subtype-aware biological comparison, future patient-stratification, biomarker, or therapeutic target-prioritization studies.

Xuanye Zhang and Anirudh Pulugurtha

Computational Characterization of Mean Arterial Pressure Trajectories After Spinal Anesthesia for Cesarean Delivery

Mentors: Prof. Vesela Kovacheva and Dr. Ricardo Kleinlein, Harvard Medical School and Brigham and Women's Hospital

Spinal anesthesia is commonly used for cesarean delivery but can cause maternal hypotension through sympathetic blockade from intrathecal bupivacaine. Phenylephrine is frequently used to prevent or treat this hypotension, yet patient responses vary, making individualized blood pressure management challenging. In this exploratory retrospective analysis, we used unsupervised clustering to identify data-driven patterns in maternal mean arterial pressure trajectories following spinal anesthesia. We analyzed intraoperative time-series data from 137 cesarean delivery cases. Trajectories were aligned to bupivacaine administration, resampled to 1-second intervals, and interpolated. Signals were then adjusted for phenylephrine administration, smoothed, and normalized to each patient's pre-spinal baseline. Dynamic time warping with k-means and three clusters produced the best separation (silhouette score 0.2478), followed by agglomerative clustering with cosine distance (0.2282). Clusters differed primarily in recovery after the blood pressure decline rather than the magnitude of the initial decline. We further evaluated early cluster prediction using a multilayer perceptron and support vector classifier, with and without TimesFM feature extraction. Across the three classes, the best model achieved 0.568 validation accuracy on 37 held-out patients (chance baseline: 0.333; binomial standard error (8%). These findings are exploratory because cluster membership may reflect baseline hemodynamics, spinal anesthesia response, clinician-directed phenylephrine administration, and preprocessing choices. This work

demonstrates the potential of unsupervised clustering to characterize heterogeneity in maternal blood pressure trajectories after spinal anesthesia. Further validation is needed to assess cluster stability, generalizability, and whether early trajectory features predict clinically meaningful hypotension patterns. Ultimately, this approach may support clinician-in-the-loop decision-making for individualized vasopressor management.

Timothy Chen and Max Levin

Autoregressive ODE Discovery from Data

Mentors: Prof. Lu Lu, Yale, and Prof. Yixiang Deng, University of Delaware

Differential equations are a natural language for many dynamical systems, concisely expressing symmetries, invariants, and fundamental principles. Given raw experimental data, recovering these governing laws can therefore offer greater interpretability and generalizability than traditional black-box methods. We introduce ODEqGPT, a transformer-based approach to this inverse problem. Building on EqGPT, which targets partial differential equations, we train a small transformer to autoregressively generate ODE token sequences, whose coefficients are separately fit via sparse regression. Candidates are then scored by cross-validated accuracy, and a fine-tuning loop reinforces top structures. We demonstrate that ODEqGPT adapts to the specific dynamics of each problem even at high observational noise. It recovers the true equations more accurately than E-SINDy, a sparse regression baseline, and its equations blow up less often when integrated. Finally, we outline BODEqGPT, which targets noisy biological dynamics with non-polynomial terms.

Alex Liu and Anthony Zou

Efficient Catalytic Algorithms for All-Pairs Shortest Paths

Mentor: Zoe Xi

Catalytic space is a model of computation in which, alongside a small amount of ordinary empty space, an algorithm is able to use a large auxiliary tape, called the catalytic tape, filled with arbitrary data – provided that it returns the filled space to its original contents by the time the algorithm halts. Recently, Cook and Pyne initiated the study of catalytic space from an algorithmic perspective, designing fast catalytic logspace algorithms for several fundamental graph problems. Continuing this research direction, in this talk we present fast catalytic logspace algorithms for the All-Pairs Shortest Paths problem, which asks to compute distances between every pair of vertices in a weighted, directed graph. Our approach yields a deterministic catalytic logspace algorithm running in time $\tilde{O}(n^{5.895})$, and a randomized algorithm that runs in time $\tilde{O}(n^{5.120})$. These improve upon the previously best-known bound of $\tilde{O}(n^{5.967})$ (for either a deterministic or randomized algorithm) [Sheffield, 2025].

Omar Graia, Caroline Huang, and Anna Zhou

Probabilistic Fault-Tolerant Metric Dimension

Mentor: Prof. Jesse Geneson, San José State University

Metric dimension asks for the minimum number of landmarks needed to uniquely identify the vertices of a graph using distances. In this talk, we study a probabilistic variant in which landmarks independently fail. Given failure probability q and confidence threshold p , we define $\dim_p(G, q)$ as the minimum deployment size whose surviving landmarks resolve the graph with probability at least p . We investigate the resulting tradeoff between deployment size and reliability across several graph families. For rectangular grids, we prove the optimality of corner configurations with up to four landmarks and derive exact resolving-subset counts and success probabilities for larger canonical boundary configurations. For trees, we characterize optimal landmark allocations using their branch structure and obtain an exact $O(N \log N)$ algorithm. We derive general bounds using correlation inequalities, binomial tail estimates, and resolving-set density. We then prove that, for dense Erdős–Rényi random graphs G , the probabilistic metric dimension satisfies $\dim_p(G, q) = \frac{1}{1-q} \dim(G) + o(\log n)$ asymptotically almost surely, for fixed $p, q \in (0, 1)$. Finally, we study the computational complexity of the problem: exact

reliability evaluation and deployment optimization are computationally hard, while the failure probability nevertheless admits an efficient randomized approximation scheme even when failures are rare. These results give structural, asymptotic, and computational perspectives on reliable distance-based identification in graphs.

SESSION 9

Jiaheng (Gary) Li

On Signature Characters and Unitary Loci for Type A_1, A_2 Quantum Groups at Roots of Unity

Mentor: Frank Wang

Quantum groups arise as q -deformations of universal enveloping algebras, playing beautiful roles in mathematical physics and integrable systems. One central problem in their representation theory is determining when a representation is unitary, which can be studied by computing signature characters that record the signature of an invariant Hermitian form on each weight space. While explicit signature character formulas for $U_q(\mathfrak{gl}_n)$ are known in the generic case, the root-of-unity case is substantially more complicated in its quantum groups structures and remains open.

In this work, we present closed-form derivations of signature characters and unitary loci for Type A_1 and A_2 quantum groups at roots of unity. Specifically, we discuss three distinct forms of Type A_1 : the De Concini-Kac quantum group, the small quantum group, and Lusztig's divided-power quantum group, while for Type A_2 , we exclusively study Lusztig's quantum group. For the rank-one De Concini-Kac quantum group, we separately discuss the highest-weight modules with dimension $< \ell$, the ℓ -dimensional fully-cyclic modules, and the ℓ -dimensional highest- and lowest-weight modules, where ℓ is the order of q^2 . For the highest- and lowest-weight modules, we also show that the density of roots of unity giving a unitary representation is zero. Finally, we present alternative proofs for the generic cases by algebraically constructing a modified Poincaré–Birkhoff–Witt basis and check that they agree with Venkateswaran's 2014 Gelfand–Tsetlin formulas.

Harry Gao

Semisimplification of the Coordinate Ring of $GL(mJ_1 \oplus nJ_{p-1})$ in Positive Characteristic

Mentor: Dr. Arun Kannan, DRW

Semisimplification provides a bridge from classical to super representation theory, turning representations of $GL(m + n(p - 1))$ into data about the supergroup $GL(m|n)$. We work in the symmetric tensor category $\mathcal{C} = \text{Rep } k[t]/(t^p)$ over an algebraically closed field of characteristic $p > 2$, and apply the semisimplification functor to the Schur algebra $S(X, d)$ with $X = mJ_1 \oplus nJ_{p-1}$. We show that $S(X, d)$ is exactly the Schur superalgebra $S(m|n, d)$ when $d < p$, and that for $d \geq p$ it is an extension of $S(m|n, d)$ by a nilpotent ideal. Dually, the extra generators of the semisimplified coordinate ring are Frobenius-twisted coordinate functions on the Jacobson radical of the centralizer of t acting on X . Since the ideal is nilpotent, the simple modules over $S(X, d)$ are exactly those inflated from $S(m|n, d)$. Finally, we extend these results to the coordinate ring $\mathcal{O}(GL(X))$, whose semisimplification is obtained from that of $\mathcal{O}(\mathfrak{gl}(X))$ by inverting the image of the determinant.

Derek Aoki

On Open Borel Orbits in Conormal Bundles to Schubert Cells

Mentor: Hamilton Wan

Let $C(v)$ denote a Schubert cell in the flag variety GL_n/B , and let $X(w)$ be a Schubert variety containing $C(v)$. Mellit has conjectured that, if $X(w)$ is smooth, the existence of an open B -orbit in the conormal bundle to $C(v)$ in $X(w)$ is equivalent to $C(w_0w)$ being contained in the smooth locus of $X(w_0v)$. We prove a stronger version of this conjecture, requiring only that $X(w)$ is factorial. We prove our result by constructing a sequence of graphs associated to the pair v, w such that the existence of a cycle in one of these graphs is equivalent to the non-existence of such an open orbit.

Jessica Yan

Computing Isomorphisms Between Groups Defined by Integer Matrices

Mentor: Prof. Maria Sabitova, CCNY – CUNY

We study the additive subgroups G_A and G_B of $\mathbb{Q}^n$ defined by nonsingular integer $n \times n$ matrices A and B , respectively. Such groups arise in the study of toroidal solenoids, $\mathbb{Z}^n$ -odometers, and symbolic dynamics. Previously, Sabitova established the computability of $\text{Iso}(G_A, G_B)$, the set of isomorphisms of abstract groups, in the 2×2 case and in the general $n \times n$ irreducible case. In this talk, we discuss the computability of $\text{Iso}(G_A, G_B)$ in the $n \times n$ reducible case as well as a general framework for constructing this set. We also examine isomorphisms between parameterized families of these rank-2 groups.

SESSION 10

Burak Aricioglu

A Closed Formula for Eigenweights

Mentor: Zeyu Wang

The eigenweights are combinatorial invariants attached to a reductive group and a dominant weight. They appear in the formula of Feng–Yun–Zhang for the arithmetic volume of the moduli stack of shtukas. Wang and Wei have described them using two bilinear forms, but for fundamental weights the resulting formulas are case-by-case and difficult to evaluate. In this talk, we give a new closed formula for eigenweights in type A; our only input is the Cartan matrix of a Dynkin diagram, which eliminates case-by-case complexity and makes them significantly easier to compute. We also explain how these methods extend beyond type A to a general formula for eigenweights across all types.

Wesley Lin

A Representation-Theoretic Computation of the Alexander Polynomial of Satellite Knots

Mentor: Kenta Suzuki, Princeton University

Let K be a knot in $\mathbb{R}^3$, and let P be a knot inside the solid torus. The satellite knot $P(K)$ is obtained by gluing P into a tubular neighborhood of K . Seifert computed the Alexander polynomial, a classic knot invariant, of $P(K)$ in terms of the Alexander polynomials of P and K . We give a novel representation-theoretic proof of this formula when the winding number of P is nonzero.

Leo Hong and Ari Shpilman

Orbital Integrals and Shalika Germs

Mentor: Dr. Thomas Rüd

Orbital integrals measure how conjugacy classes of matrices meet prescribed arithmetic conditions, such as having integral coefficients. For p -adic groups, their behavior near zero can be described by *Shalika germs*, which are coefficients associated with each nilpotent conjugacy class. We describe how to calculate these germs in the case of elliptic equivalued points of $\mathfrak{sl}_n$ using specific test functions and a Cartan decomposition which reduces the computation to point counts on certain flag varieties.

This geometric viewpoint provides an explicit case study within an established numerical theory. It also prepares a concrete next question: in Hales’s equivalued example in $\mathfrak{sp}_6$, where a stable subregular germ contains an elliptic curve point count, do other nilpotent germs exhibit comparable arithmetic behaviors? Another broader motivation comes from the proposed primitive trace formula and the need to understand the local contributions relevant to that program.

Brian Sun and Seabert Mao

Smoothness Characterizations and Duality of the Type II Theta Correspondence and Basic Fourier–Jacobi Models

Mentor: Dr. Jialiang Zou

We study the smoothness of two families of Borel-orbit closures. The first consists of matrix Schubert varieties in $\text{Mat}_{n \times m}$. The second consists of orbit closures for the diagonal action of the upper-triangular Borel subgroup $B \subseteq \text{GL}(V)$ on $\text{FL}(V) \times V$, where $V = \mathbb{C}^n$.

For matrix Schubert varieties, we prove that smoothness is equivalent to the absence of a contiguous 2×2 block of equal positive entries in the southwest rank table of the indexing partial permutation matrix. The proof also shows that every smooth matrix Schubert variety is a coordinate linear subspace. We enumerate these varieties both in total and by rank; the rank-refined counts are Catalan triangle numbers.

The B -orbits in $\text{FL}(V) \times V$ are indexed by colored permutations. We give a complete smoothness criterion involving two ordinary permutation patterns and two colored patterns, together with two exceptional extensions of the final colored pattern. For fixed n , we enumerate the smooth orbit closures in this setting, obtaining a Catalan-like recurrence involving the number of smooth Schubert varieties of a fixed dimension.

When $m = n$, the Borel orbits in the two ambient spaces are naturally in bijection, and hence so are their closures. In the framework of relative Langlands duality, this is the expected correspondence between the type II theta-correspondence model and the basic Fourier–Jacobi model. We show that the bijection does not preserve smoothness in general and characterize the failures in one direction.

Grant Blitz

Coprimality and Divisibility in Integer Matrices

Mentor: Prof. Maxim Arnold, University of Texas Dallas

Given two integer matrices, we establish a notion of coprimality, using an analog of Bezout’s Lemma over the integers. We explore the consequences of this condition, providing a classification of all matrices that are coprime under this notion. Then we turn to investigating common divisors of matrices, defining the greatest common divisor of two matrices to be a specific basis of the lattice generated by their column vectors. We explore the consequences of this definition, and show that many properties of greatest common divisors from the integers have analogs in matrices. Last, we introduce and investigate prime-like matrices, showing they inherit some properties of primes in the integers.

Niranjan Rao

Cubic Surfaces over $\mathbb{F}_q$ with 27 Lines over $\mathbb{F}_{q^2}$

Mentor: Xinyu Fang, Harvard University

We study cubic surfaces over $\mathbb{F}_q$ that contain 27 lines defined over $\mathbb{F}_{q^2}$. We show all such cubic surfaces have at least one real tritangent plane. We propose a procedure to construct a class of these cubic surfaces, adapting a method used by Betten and Karaoglu (2022) to compute the number of cubic surfaces over $\mathbb{F}_q$ with 27 rational lines. We also propose a method to compute the explicit equations of the 27 lines on cubic surfaces in this class, as well as the equations of the tritangent planes over $\mathbb{F}_q$, building on the arc-lifting algorithm developed by Betten, Hirschfeld and Karaoglu (2018).

Abhigyan Singh

T-equivariant Modules over the Ring $k[x_1, x_2, \dots, x_n][y_1, y_2, \dots, y_{n-1}]/(x_i y_i)_{i=1}^{n-1}$

Mentor: Dr. Linus Hamann, Harvard

Homological algebra of modules is used to studying the failure of certain sequences to be *exact*, in particular allowing for tools that analyze the “number” of ways objects interact in higher degree extensions. Studying these modules equipped with a $\mathcal{O}(T)$ -coaction, for T a torus, considers this same behavior, but investigates it when the maps are forced to be equivariant with respect to this action. In this talk, we introduce some machinery used in homological algebra, in particular covering the notion of a projective module and projective resolutions. We then introduce, for the case when $R = k[x_1, x_2, y_1]/(x_1 y_1)$, a torus action on the geometric space $\text{Spec}(R)$ associated to R , and then state an equivalent interpretation from the perspective of modules, covering the concept of a T -equivariant module. After discussing this setup, we briefly note that these objects can be organized into a *derived category*, and present some results describing various projective resolutions and higher degree extensions in this category.

SESSION 12

Madeline Hu

The Associative Zesting of an Exact Factorization of Fusion Categories

Mentors: Prof. Julia Plavnik and Dr. Héctor Peña Pollastri, Indiana University Bloomington

Fusion categories are mathematical objects with major implications for topological quantum field theory, causing methods of constructing new fusion categories to be of utmost importance. Two such constructions are associative zesting, which uses cohomology to modify the structure of an existing category, and exact factorization, which combines two existing categories.

In this presentation, we explore how these two constructions intertwine when fusion categories are graded by their universal grading groups. We begin by developing a notion of an “induced” zesting from a fusion category on a fusion subcategory when the zesting satisfies a stabilization condition. With this notion of an induced zesting, we show that the induced zestings of an exact factorization preserve an exact factorization of the parent zesting. We develop a second, more general method of inducing a zesting from a parent fusion category. We prove that these induced zestings are entirely determined by factorizations of characters and a factorization of 3-cocycles. Moreover, we show that these induced zestings also preserve an exact factorization of fusion categories, indicating compatibility between the two constructions.

Tarun Rapaka

Unfactorizable Zestings of Deligne Products of Tambara–Yamagami Fusion Categories

Mentors: Prof. Julia Plavnik and Dr. Héctor Peña Pollastri, Indiana University Bloomington

Fusion categories are algebraic structures that arise in several areas of mathematics and mathematical physics. An important problem in their study is the construction of new examples and the understanding of their properties. In this talk, we will show how associative zesting can be used to construct new fusion categories from the Deligne product of two Tambara–Yamagami categories. We will examine the structures of these new categories, focusing on their factorization properties. Finally, we will turn our attention to the special case of two Ising categories.

Dannan Li

A Bound on the Fractal Uncertainty Principle for a Square Alphabet

Mentor: Alain Kangabire

We prove a lower bound of $\frac{1}{2} - \frac{7\delta}{8} - \epsilon$ ($0 < \delta \leq \frac{1}{4}$) on the fractal uncertainty exponent β for dilated Cantor sets $C_k(N)$ when the alphabet $\mathcal{A}$ is given by squares, thus leading to an improvement of $\frac{\delta}{8}$ on the general bound of $\frac{1}{2} - \delta$ for undilated Cantor sets. If we restrict the dilation parameter to rational numbers with big denominator, we prove a stronger lower bound of $\frac{1}{2} - \frac{3\delta}{4} - \epsilon$ on the fractal uncertainty exponent, which gives an improvement of $\frac{\delta}{4}$. As an application, these improved fractal uncertainty exponents yield better spectral gaps for open quantum Baker's maps.

SESSION 13

Charlie Gu and Krittika Chandra

Relative Humidity and Wind Speed for Long-Term Probabilistic Load Forecasting

Mentor: Dr. Tao Hong, University of North Carolina at Charlotte

Electric load forecasting is the process of predicting future electricity demand through historical data and explanatory variables. The most widely used explanatory weather variable is temperature. Others, such as wind speed and relative humidity, have also been studied in the load forecasting literature, though they are rarely adopted in practice. We recreate two high-performing linear regression models from the past decade, one with wind speed terms and the other with relative humidity terms, and compare their performance to Tao's Vanilla Benchmark Model and its extension with recency effect temperature terms. We then propose a new comprehensive model including both wind speed and relative humidity terms. These models are tested for year-ahead point forecasting and then extended to year-ahead probabilistic forecasting. The case study is performed on Boston load data from ISO New England, and weather conditions from Boston Logan International Airport. For point forecasting, the wind speed model and relative humidity model are both shown to outperform the Vanilla Benchmark and recency effect models in all cases, and the comprehensive model is shown to outperform all four prior models in all cases, demonstrating the value of including both weather variables simultaneously. The same is found for long-term probabilistic forecasting, across all scenario generation settings explored. In particular, depending on the scenario generation method, the comprehensive model demonstrates a 1.53%-3.14% improvement in quantile score compared to the Vanilla Benchmark and a 0.63%-1.45% improvement in quantile score compared to the recency effect model. This work demonstrates the value of relative humidity and wind speed variables even in probabilistic forecasting where their exact values are unknown.

Aarush Kulkarni, Sirui (Kris) Sheng, and James Wu

Directional Curvature for Efficient Word2Vec Optimization

Mentor: Dr. Willy Rodriguez, Torus AI

Word embeddings are a foundation of many natural language processing applications and they are usually trained using first order methods, which use only gradient information, such as Adam or SGD. In this talk we discuss how second order curvature information, from the matrix of second derivatives, can be used to improve the optimization of the objectives for both the CBOW and skip gram in word2vec, two models that learn embeddings by predicting words from their neighbors in text. We present a closed formula for its directional curvature, the second derivative along a given direction, and show that it can also be computed at the same time as the gradient. We compare our implementations of the Rescaled Exponential Decay (RED) method, two new extensions of it (RED-M and RED-Hessian), and several

other baseline methods. Our experiments show that curvature-rescaled step sizes can give a more stable convergence behavior, reaching a loss comparable to well-tuned first-order baselines. However, lower training and test loss did not consistently translate into higher-quality semantic embeddings, especially for the optimizers that apply curvature information to the search direction.

Yunong (Melody) Wu, Gary Shen, and Tanish Parida

Predicting Cognitive Decline from Baseline Assessments: An OASIS-3 Study

Mentor: Dr. Marly Gotti, Arcus Biosciences

Cognitive decline, including that associated with Alzheimer’s disease, is a major global concern that often goes unrecognized until it begins to interfere with daily life. We investigated whether routinely available information from an older adult’s first clinical visit could identify individuals at risk of decline over the next two to four years. Using longitudinal data from 458 adults aged 60 and older in the OASIS-3 cohort, we compared logistic regression, random forest, and gradient-boosting models based on demographic, clinical, and cognitive measures. All models outperformed chance, with the clinical-and-cognitive logistic regression model achieving the best performance, indicating baseline clinical status as especially informative. These results establish a practical non-imaging benchmark for evaluating whether more complex MRI-based models offer meaningful additional predictive value to cognitive decline prediction.