Union College Math Conference: Computational Mathematics

SEPTEMBER 18–20, 2026

Friday program

4:30–5:30pm: Reception and registration (Olin Rotunda)

5:30–6:30pm: Plenary talk I: Carina Curto: *Network Dominoes: Composable Building Blocks for Recurrent Networks* (Olin 115)

Saturday program

8:00–9:00am: Registration, coffee & light breakfast (Bailey 204)

8:30–10:30am: Session I (ISEC 222)

- 8:30–9:00 Sam Hawke: *Balancing fairness and structure in dimension reduction*
- 9:00–9:30 Minjie Wang: *High-Dimensional Variable Selection with Diffusion-Generated Synthetic Data*
- 9:30–10:00 Halyun Jeong: *Kaczmarz and one-pass methods for sparse recovery and robust streaming regression*
- 10:00–10:30 Andrew Horning: *High-order Matrix-Free Spectral Density Estimation*

10:30–11:00am: Coffee break (Olin Rotunda)

11:00am–12:00pm: Plenary talk II: Kristine Bauer: *The Categorical Derivative in Functor Calculus* (Olin 115)

12:00–1:30pm: Lunch break

1:30–3:30pm: Session II (ISEC 222)

- 1:30–2:00 Yoonsang Lee: *Stochastic reduced-order modeling for fluid systems*
- 2:00–2:30 Nicholas J. Moore: *Reversals of the large-scale circulation in thermal convection*
- 2:30–3:00 Sayantan Sarkar: *From video to equations: what a discovered coefficient actually measures*

- 3:00–3:30 Ziyao Xu: *TVD and TVB Preservation in RKDG and cRKDG Methods without TVD Time Discretization*

3:30–4:00pm: Coffee break (Olin Rotunda)

4:00–5:00pm: Plenary talk III: Julie Blackwood: *Nonlinear Leslie matrix models across scales: from the evolution of bird song to cicada emergence* (Olin 115)

5:00–6:00 pm: Reception (Olin Rotunda)

Sunday program

8:00–8:30am: Coffee & light breakfast (Bailey 204)

8:30–10:30am: Session III (ISEC 222)

- 8:30–9:00 Shukai Du: *Accelerating finite element via machine learning on one element*
- 9:00–9:30 Emmanuel Adara: *Beyond traditional solvers: MLPs and KANs for the Chemical Master Equation*
- 9:30–10:00 Jihun Han: *A stochastic vanishing viscosity approach for eikonal equations*
- 10:00–10:30 Andrea Dziubek: *Advanced Numerical Methods for Materials and Manufacturing*

10:30–11:00am: Coffee break (Olin Rotunda)

11:00am–12:00pm: Plenary talk: Lan-Hsuan Huang: *Bartnik's Quasi-Local Mass in General Relativity* (Olin 115)

12:00–1:30pm: Lunch break

1:30–3:30pm: Session IV (ISEC 222)

- 1:30–2:00 Felix Ye: *FlashSinkhorn: IO-Aware Entropic Optimal Transport on GPU*
- 2:00–2:30 Aimee Maurais: *Localization of flow-based samplers and generative models*
- 2:30–3:00 Yangyang Xu: *Scalable Dynamic Optimal Transport via Distributed Linearized ADMM*
- 3:00–3:30 Sean Carney: *Optimistic optimal transport for inverse data matching with corrupted data*

Abstracts

Emmanuel Adara (State University of New York at Binghamton)
Beyond traditional solvers: MLPs and KANs for the Chemical Master Equation
Accurately modeling the stochastic dynamics of chemical reaction networks is fundamental to understanding and predicting the behavior of complex biochemical systems. The Chemical Master Equation (CME) provides a rigorous probabilistic framework for describing the time evolution of molecular populations. However, solving the CME remains a significant computational challenge due to the *curse of dimensionality*, as the size of the state space grows exponentially with the number of chemical species, making conventional numerical methods computationally prohibitive for all but relatively small systems.
In this talk, I will present an overview of classical computational approaches for analyzing and approximating the CME, highlighting their strengths and limitations. I will also discuss recent advances in scientific machine learning, focusing on neural surrogate models based on Multi-Layer Perceptrons (MLPs) and Kolmogorov–Arnold Networks (KANs). These data-driven approaches provide efficient and scalable approximations of CME solutions, offering a promising framework for accelerating the analysis and simulation of high-dimensional stochastic reaction networks.

Julie Blackwood (Williams College)
Plenary talk: *Nonlinear Leslie matrix models across scales: from the evolution of bird song to cicada emergence*
Leslie matrix models are commonly used to capture the population dynamics of structured populations. This talk explores the versatility of this framework across multiple scales: the cultural evolution of bird song (which occurs on a relatively fast time scale for evolution) and the spatial synchrony of periodical cicadas (across long ecological time scales and large spatial scales). First, we examine a multi-decade cultural shift in an island population of Savannah sparrows, where "high note clusters" were entirely replaced and elaborated on by "click trains." By applying a structured modeling framework to track song transmission, we show that this naturally occurring cultural evolution was likely driven by direct selection, rather than drift or frequency-dependent selection. Second, we investigate the 17-year emergence cycles of periodical cicadas in the U.S. Using a nonlinear Leslie-type matrix model, we show how the relationship between nymphal competition and the ability of some individuals to delay or accelerate their emergence as adults can lead to large-scale synchrony.

Sean Carney (Union College)
Optimistic optimal transport for inverse data matching with corrupted data
Data normalization is a common issue that arises when applying the quadratic Wasserstein metric from optimal transport (OT) theory to inverse problems with signed signals. The raw signals must be transformed into non-negative functions with equal mass for OT theory to be applicable. This work considers an "optimistic optimization" approach to this issue based on problem relaxation. Instead of normalizing signals with a single, fixed transformation, this approach dynamically normalizes the data signal based on a variational principle. This leads naturally to a misfit function that can be used in the numerical solution of inverse problems. We describe the mathematical properties of the misfit function, including results that show its robustness to noisy data corrupted with outliers. Using synthetic examples from seismic imaging problems, we demonstrate how the approach can handle corrupted data that cause issues for existing methods.

Shukai Du (Syracuse University)
Accelerating finite element via machine learning on one element
Machine learning (ML) has shown potential for accelerating numerical simulations, with notable examples such as numerical weather prediction and molecular dynamics using ML-predicted potentials. At the same time, many data-driven ML approaches face well-known challenges, such as limited interpretability and reliability, as well as substantial demands for training data and computational resources.
I will talk about a new ML-based approach to accelerate classical finite element methods while preserving their key advantages, including interpretability, robustness, and applicability to complex geometries. The central idea is to train an ML model on a single reference element, making training inexpensive and reducing data requirements. Once trained, the model is used to accelerate computation on the whole domain with the techniques of hybridization and static condensation. I will demonstrate the approach on the radiative transfer problem and report numerical results. If time permits, I will discuss ongoing work of the theoretical analysis of the method.

Andrea Dziubek (State University of New York Polytechnic Institute)
Advanced Numerical Methods for Materials and Manufacturing
The mathematical methods required for modern computational mechanics and advanced manufacturing have matured over the last several decades. They are Geometric Mechanics and Structure-Preserving Numerical Methods for ODEs, PDEs, and Optimal Control Theory.
Geometric Mechanics is Continuum Mechanics on manifolds, in the language of differential geometry and exterior calculus – essential for modeling complex materials such as Cosserat media with rotational degrees of freedom (values in a vector bundle, not in a vector space).
Symplectic integrators preserve the Hamiltonian structure of conservative systems (celestial mechanics, molecular dynamics). Finite Element Exterior Calculus (FEEC) preserves the geometric and topological structure of PDEs through the de Rham complex (Maxwell's equations, elasticity, and fluid mechanics). Metriplectic integrators preserve energy and entropy for dissipative systems (thermoelasticity, magnetohydrodynamics).
Optimal Control Theory and Deep Learning share the same mathematical structure, suggesting that structure-preserving methods are the foundation for reliable AI in scientific computing.
Research in these fields is often only accessible at top-tier universities, and these topics have yet to make it into standard curricula. It might be time. *Knowledge needs to be nourished to flourish.*

Jihun Han (State University of New York at Albany)
A stochastic vanishing viscosity approach for eikonal equations
We present a meshless stochastic vanishing viscosity approach for solving eikonal equations. By introducing a small viscosity, we derive a stochastic formulation that captures directional transport along gradient trajectories. A neural network is trained to learn this transport structure rather than enforcing pointwise PDE residuals. By appropriately scaling the stochastic timestep, the method maintains a consistent spatial displacement, leading to viscosity-independent computational cost without iterative viscosity reduction. We demonstrate the method for various source configurations, anisotropic and spatially varying speeds, and non-convex domains with internal obstacles. By adapting the stochastic process to the local geometry near obstacles, the network learns appropriate boundary behavior without explicitly imposing additional constraints such as the Soner constraint.

Sam Hawke (Skidmore College)
Balancing fairness and structure in dimension reduction
Dimension reduction (DR) helps researchers analyze high-dimensional data by replacing many variables with fewer informative directions. But what counts as "informative"? In some settings, we want to preserve features most useful for predicting an outcome; this is known as sufficient DR. In biological applications such as case-control studies, we may instead wish to isolate structure that is unique to a target group relative to a background group; this is called contrastive DR. Fair DR introduces a different challenge: constructing a low-dimensional representation that removes information about a sensitive attribute. In this talk, I will introduce the FairRe framework, which addresses this challenge by balancing three goals: predictive performance, fairness, and preservation of the original data structure. FairRe draws on mathematical connections with sufficient and contrastive DR. I will describe the main ideas behind the framework, illustrate its performance through real-data experiments, and discuss possible directions for future work.

Andrew Horning (Rensselaer Polytechnic Institute)
High-order Matrix-Free Spectral Density Estimation
We introduce a new matrix-free algorithm for computing spectral densities of large "mat-vec" matrices, and use it to compute the local density of states (LDOS) of sparse Hamiltonians derived from tight-binding models. Our approach can be viewed as a high-order variant of the popular regularized Chebyshev kernel polynomial method (KPM), which uses a high-order accurate approximation of the δ-function to achieve rapid convergence to the thermodynamic limit for smooth spectral densities. The computationally intensive steps are identical to those for KPM, and high-order accuracy is achieved with an inexpensive post-processing procedure. We apply the algorithm to tight-binding models of graphene and twisted bilayer graphene, demonstrating high-order convergence to the LDOS at non-singular points.

Halyun Jeong (State University of New York at Albany)
Kaczmarz and one-pass methods for sparse recovery and robust streaming regression
This talk presents two recent directions in simple iterative methods for regression problems involving structural constraints and corrupted observations. First, I will discuss reshuffling Kaczmarz methods with sparse constraints for linear inverse problems, where Kaczmarz sweeps combined with hard thresholding lead to efficient algorithms with provable linear convergence under standard random sensing models. I will then turn to strict one-pass methods for linear and ReLU regression under response corruptions, where a sign-based stochastic method with exponentially decaying step sizes yields robust guarantees suited to streaming settings. In this setting, a sign-based stochastic method with exponentially decaying step sizes provides robust guarantees tailored to streaming data. The unifying theme is that simple iterative updates can accommodate structural constraints, robustness to corruption, and limited access to observations. This talk is based on joint work with Deanna Needell and Elizaveta Rebrova.

Yoonsang Lee (Dartmouth College)
Stochastic reduced-order modeling for fluid systems
This talk discusses several stochastic approaches—such as stochastic superparameterization and conditional Gaussian closure—for multiscale reduced-order modeling of fluid systems, focusing on stability, accuracy, and efficiency. We particularly focus on the Feynman–Kac formulation of the viscous incompressible Navier–Stokes equations to introduce a derivative-free method. This method is mesh-free, avoiding the complicated coarse-grid design typically

required to capture macroscopic behavior. We show that the method captures the large-scale dynamics of turbulent systems at a fraction of the computational cost of standard methods.

Aimee Maurais (Cornell University)
Localization of flow-based samplers and generative models
Sampling from a probability distribution is fundamental to modern computational mathematics. A powerful, emerging approach to sampling is dynamic measure transport (DMT): the idea is to design dynamics that evolve samples from a tractable reference distribution to the desired target distribution. DMT approaches for sampling generally require identification of a suitable velocity field or drift, which can be challenging in high dimensions. In this talk I demonstrate how sparse conditional dependence structure can be exploited to enhance scalability of DMT approaches for sampling. I show that commonly used velocities for DMT may inherit locality structure from the target distribution in the sense that they can be well-approximated by velocities with sparse variable dependence. I show that the use of localized parametrizations to learn DMT velocities may lead to favorable bias-variance tradeoffs when data available for training are scarce and thus lead to better sampling performance. This talk is based on joint work with Shuigen Liu (Penn State), Youssef Marzouk (MIT), and Xin Tong (NUS).

Nicholas J. Moore (Colgate University)
Reversals of the large-scale circulation in thermal convection
The large-scale circulation (LSC) associated with thermal convection is known to spontaneously reverse direction. In the atmosphere, reversals can result in a sudden change in wind direction, while in the liquid core, reversals may play a role in magnetic dipole shifts. We examine LSC reversals within the context of thermal convection in an annular domain. Through comparison with direct numerical simulations, we show that a low-dimensional dynamical system derived systematically from Galerkin truncation of the governing equations accurately describes a sequence of parameter bifurcations, including the onset of circulatory flow, the appearance of chaotic LSC reversals, and finally a high-Rayleigh-number state of periodic LSC reversals with small-scale turbulence. When cast in terms of the fluid's angular momentum and center of mass, the model reveals equivalence to a pendulum system with driving term that raises the center of mass above the fulcrum. It is the competition between driving, restoring, and damping that leads to the range of convective states. This physical picture yields accurate predictions for the frequency of regular LSC reversals in the high Rayleigh-number limit and offers a transparent mechanism for reversals.

Sayantan Sarkar (University of Massachusetts Amherst)
From video to equations: what a discovered coefficient actually measures
We present a staged pipeline that recovers an explicit, simulable partial differential equation from uncalibrated video of a spreading dye droplet: preprocessing, an optical-density transformation, centroid drift estimation, weak-form sparse regression over candidate libraries, and forward validation on held-out frames. It recovers prescribed coefficients from synthetic data to within three percent and predicts held-out experimental frames to within 3.5 percent, selecting a model dominated by a large nonlinear gradient term. We then ask what that term measures. Treating the unknown camera response as a change of state variable, its coefficient transforms as a connection on the state space, and in one state dimension every such connection is flat: the coefficient takes any value one likes under a change of observation variable.

Read through this lens, the fitted numbers return an implausible and non-unique instrument calibration curve rather than transport physics.

Minjie Wang (The State University of New York at Binghamton)
High-Dimensional Variable Selection with Diffusion-Generated Synthetic Data
Recent advances in generative models and AI have catalyzed synthetic data generation, which challenges traditional statistical practice of focusing on raw data and shifts toward a synthetic data-centric approach. Meanwhile, variable selection for high-dimensional, highly correlated data has long been a challenging problem. To enhance variable selection, we present a resample-aggregate framework that exploits diffusion models' ability to generate high-fidelity synthetic data. Specifically, we draw multiple pseudo-data sets from a diffusion model fitted to the original data, apply any off-the-slate selector (e.g., lasso), and store the resulting inclusion indicators. Aggregating across replicas produces a stable subset of predictors with calibrated stability scores for variable selection. Theoretically, we show the proposed method is selection consistent under mild assumptions. Extensive simulations show consistent gains over existing methods, especially when predictors are strongly correlated. By coupling diffusion-based data augmentation with principled aggregation, our method advances variable selection for interpretable, statistically rigorous scientific analysis.

Yangyang Xu (Rensselaer Polytechnic Institute)
Scalable Dynamic Optimal Transport via Distributed Linearized ADMM
In this talk, we address two fundamental challenges in the numerical solution of dynamic optimal transport (OT) problems. The first challenge arises when the initial and/or terminal densities approach zero and no positive lower bound is available. In this regime, conventional methods may become unstable or computationally inefficient, since the gradient Lipschitz constant of the objective can scale like the reciprocal of the cube of the density. As a result, near-zero regions may lead to slow convergence or even divergence. The second challenge concerns the substantial memory cost of the dynamic formulation, whose discretization over fine spatial and temporal grids requires storing variables across the entire space-time domain. This storage burden quickly becomes prohibitive as the grid is refined or the spatial dimension increases. To overcome the first difficulty, we reformulate the classical discretized dynamic OT problem so that the resulting objective admits an exact proximal mapping. When combined with a linearized alternating direction method of multipliers (LADMM), this reformulation yields an efficient and robust algorithm that remains stable even in challenging settings where the density may vanish. To reduce the memory burden, we further introduce a distributed formulation in which the optimization variables are partitioned across multiple agents. This design substantially lowers the storage requirement for each agent and can also lead to computational acceleration. We validate the proposed framework through numerical experiments in one- and two-dimensional spatial settings under varying levels of difficulty. The results demonstrate the stability, robustness, and scalability of the proposed method in comparison with conventional approaches.

Ziyao Xu (State University of New York at Binghamton)
TVD and TVB Preservation in RKDG and cRKDG Methods without TVD Time Discretization
Total variation diminishing (TVD) and total variation bounded (TVB) properties are essential for controlling spurious oscillations in numerical conservation laws. In the classical

Runge–Kutta (RK) discontinuous Galerkin (DG) framework, enforcing TVD or TVB properties is intrinsically tied to TVD time integrators, which are more commonly known today as strong-stability-preserving (SSP) methods. This reliance imposes severe structural restrictions, including order barriers and a lack of compatibility with the recently developed RKDG method with compact stencil (cRKDG). To bypass these constraints, we propose a novel trace-limited corrector framework that preserves the TVD/TVB-in-the-means properties using generic, non-SSP RK time discretizations. Based on Harten's lemma, our key insight is that only the cell-average update in the final corrector stage is essential for enforcing total variation stability. Consequently, we modify the traces in the numerical fluxes exclusively in the final stage, leaving the intermediate predictor stages unconstrained. This strategy accommodates both standard RKDG and cRKDG predictors while maintaining strict conservation and the classical TVD/TVB interpretations. By eliminating the need for immediate limiting operations after each intermediate stage, the proposed method decouples oscillation control from the SSP restriction, retains the compactness of the cRKDG framework, and offers a streamlined algorithmic flow. We also show that the limiting procedure preserves the underlying accuracy of the scheme by proving its non-activation property in smooth regions. Finally, numerical experiments are presented to demonstrate the capabilities and robustness of the method.

Felix Ye (State University of New York at Albany)
FlashSinkhorn: IO-Aware Entropic Optimal Transport on GPU
Entropic optimal transport (EOT) via Sinkhorn iterations is widely used in modern machine learning, yet GPU solvers remain inefficient at scale. Tensorized implementations suffer quadratic HBM traffic from dense $n \times m$ interactions, while existing online backends avoid storing dense matrices but still rely on generic tiled map-reduce reduction kernels with limited fusion. We present *FlashSinkhorn*, an IO-aware EOT solver for squared Euclidean cost that rewrites stabilized log-domain Sinkhorn updates as row-wise LogSumExp reductions of biased dot-product scores, the same normalization as transformer attention. This enables FlashAttention-style fusion and tiling: fused Triton kernels stream tiles through on-chip SRAM and update dual potentials in a single pass, substantially reducing HBM IO per iteration while retaining linear-memory operations. We further provide streaming kernels for transport application, enabling scalable first- and second-order optimization. On A100 GPUs, FlashSinkhorn achieves up to 32× forward-pass and 161× end-to-end speedups over state-of-the-art online baselines on point-cloud OT, improves scalability on OT-based downstream tasks.