ACM 2026

Applied and Computational Mathematics Workshop Program Booklet

April 3–4, 2026

Plenary Talks
Student Talks
Posters

Department of Mathematics
Emory University
Atlanta, GA

Contents

Program Overview **2**
- Friday, April 3 . . . 2
- Saturday, April 4 . . . 3

Plenary Talks **5**

Student Talks **6**
- Student Talk Session I . . . 6
- Student Talk Session II . . . 9
- Student Talk Session III . . . 11
- Student Talk Session IV . . . 13
- Student Talk Session V . . . 15

Posters **17**

Program Overview

Friday, April 3

Friday Schedule	12:30 PM – 6:15 PM
12:30 – 1:00 PM	Check-in / Registration
1:00 – 1:15 PM	Opening Remarks Jim Nagy
1:15 – 2:00 PM	**Plenary Talk 1** — Richard Lehoucq (Sandia National Laboratories) *The Poisson Tensor Completion Parametric Estimator*
2:00 – 3:30 PM	**Student Talk Session I** Mitchell Scott — What Makes a Good Data Science Preconditioner? (Emory University) Jonathan Valyou — A Conjugate-Gradient Formulation of the EnKF Algorithm (Florida State University) Atticus Rex — Multifidelity-Augmented Gaussian Process Inputs for Surrogate Modeling from Scarce Data (Georgia Tech) Francesco Brarda — Preconditioning Via Spectral Density Driven Graph Neural Networks (Emory University) Haoyang Qian — When Complexity Simplifies: The Paradox of More Random, More Tractable (Florida State University)
3:30 – 4:00 PM	Coffee Break
4:00 – 4:45 PM	**Plenary Talk 2** — Lars Ruthotto (Emory University) *Optimal Transport and PDE Perspectives on Generative AI*
4:45 – 6:15 PM	**Student Talk Session II** Pengjun Wang — Error Estimates of a Training-Free Diffusion Model for High-Dimensional Sampling (Auburn University) Caleb Fikes — PTYDIFF: Score-Induced Implicit Phase Choice Sampling for Nonlinear Ptychographic Imaging Problems (Emory University) Jonathan Engle — Elections and the Spatial Public Goods Game (Florida State University) Munawar Ali — Noise Estimation of SDE from a Single Data Trajectory (Florida State University) Farhana Taiyebah — Accelerating Posterior Inference from Pulsar Light Curves via Learned Latent Representations and Local Simulator-Guided Optimization (Florida State University)
6:15 PM	Dinner / Poster Session (16 posters) Kapil Chawla, Srijon Sarkar, Ngoc Hien Tran, Yahong Yang, Vivian Zhang, Shikhar Shah, Anand Natarajan Sriram, Benjamin Burns, Meixi Li, Tianshi Xu, Toma Debnath, Zhongjie Shi, Callie Reid, Shiyi Lyu, Zixiang Xiong, Yellin Benjamin

Saturday, April 4

Saturday Schedule	8:00 AM – 5:15 PM
8:00 – 9:00 AM	Breakfast (provided)
9:00 – 9:45 AM	**Plenary Talk 3** — Alessandro Veneziani (Emory University) *The role of Mathematics in the Clinics of Cardiovascular Diseases: Challenges and Perspectives*
9:45 – 11:15 AM	**Student Talk Session III** Jingqiao Tang — Data Assimilation Framework for Uncertainty Reduction in Learning Data-Driven Dynamical Systems (Florida State University) Ruoyu Hu — Introduction of Power Grid with a Numerical Example on Feedback Control Problem with Data Assimilation (Florida State University) Yi Liu — Multi-level Machine Learning Framework for Inverse Scattering Problems (Auburn University) Tomoki Koike — Sparse POD Mode Selection and Manifold Dimensionality Reduction with Neural Networks (Georgia Tech) Seungmin Lee — An optimization-based bound-preserving limiter for two-phase systems (Florida State University)
11:15 – 11:45 AM	Coffee Break
11:45 AM – 12:30 PM	Panel Discussion: Careers and Research Opportunities in Scientific Computing and Data Science
12:30 – 1:45 PM	Lunch (provided)
1:45 – 3:15 PM	**Student Talk Session IV** Ferhat Karabatman — Geometric Perspective on Concentration Phenomena in Frame Theory (Florida State University) Xu (Melissa) Wang — A Primal-Dual Price-Optimization Method for Computing Equilibrium Prices in Mean-Field Games Models (Emory University) Mujtaba Ali — On the Hausdorff Stability of Barcodes over Posets (Florida State University) Jiaqi Yang — Multi-patient Computational Analysis of Type B Aortic Dissections (Emory University) Haoran Yan — Understanding Denoising Autoencoders through the Manifold Hypothesis: A Geometric Perspective (Georgia Tech)
3:15 – 3:45 PM	Coffee Break

Saturday Schedule	8:00 AM – 5:15 PM
3:45 – 5:15 PM	**Student Talk Session V** Yi-Yung Yang — Bound-preserving and Entropy Stable Enriched Galerkin Methods for Nonlinear Hyperbolic Equations (Florida State University) Kha Doan — Dynamically Regularized Lagrange Multiplier Method for the Cahn-Hilliard–Navier–Stokes System (Auburn University) Leonardo Molinari — Computational Framework for Cardiac Radiofrequency Ablation: Multiphysics Modeling and Domain Decomposition (Emory University) Akshita Sahni — A Reduced Order Model for Rapid 3D Fluid-Structure Interaction Simulation of Aortic Valve Flows (Georgia Tech) James Orgeron — Habitat Fragmentation Promotes Spatial Scale Separation Under Resource Competition (Florida State University)
6:00 PM	Closing Remarks

Plenary Talks

Richard Lehoucq Sandia National Laboratories

The Poisson Tensor Completion Parametric Estimator

We introduce the Poisson tensor completion (PTC) estimator that exploits inter-sample relationships to compute a low-rank Poisson tensor decomposition of the frequency histogram for samples of a multivariate distribution. Our crucial observation is that the histogram bins are an instance of a space partitioning of counts and thus can be identified with a spatial non-homogeneous Poisson process. The Poisson tensor decomposition leads to a completion of the mean measure over all bins, including those containing few to no samples, and leads to our proposed estimator. A Poisson tensor decomposition models the underlying distribution of the count data and guarantees non-negative estimated values, obviating the need for additional constraints to ensure non-negativity. Furthermore, we demonstrate that our PTC estimator is a substantial improvement over standard histogram-based estimators for sub-Gaussian probability distributions because of the concentration of norm phenomenon.

Lars Ruthotto Emory University

Optimal Transport and PDE Perspectives on Generative AI

Modern generative AI systems, from Stable Diffusion to Sora, produce strikingly realistic images and videos. Perhaps surprisingly, the mathematical foundations of these models rest on classical PDEs well known in the applied mathematics community.

In this talk, I will show how two fundamental PDEs, the continuity equation and the Fokker–Planck equation, provide a framework for understanding today's approaches to generative modeling. Starting from the continuity equation and the method of characteristics, I will derive continuous normalizing flows as neural ODEs and explain how the Benamou–Brenier formulation of optimal transport provides structure through energy minimization. I will then show how flow matching sidesteps costly time integration during training by constructing feasible probability paths via conditional transport between point pairs: a construction enabled by the linearity of the continuity equation. Turning to stochastic transport, I will demonstrate how the Fokker–Planck equation naturally yields the score function, and how score matching reduces to supervised learning with analytically known conditional scores.

A central theme is the distinction between feasible and optimal transport: while optimal transport provides an elegant theoretical structure, modern state-of-the-art methods succeed by constructing feasible solutions that are computationally tractable in high dimensions. I will conclude by outlining open research challenges, including the pursuit of efficient matching-type algorithms for optimal transport and connections to optimal control and mean-field games.

Alessandro Veneziani Emory University

The role of Mathematics in the Clinics of Cardiovascular Diseases: Challenges and

Perspectives

This talk is motivated by the conviction that mathematics is essential to societal progress, and that its impact on healthcare, particularly cardiovascular medicine, has yet to be fully realized. While technological transfer is often celebrated as the primary driver of innovation in medicine, the equally critical role of methodological transfer is frequently overlooked. Advances in clinical understanding and decision-making are ultimately grounded in models, methods, and analytical frameworks that rely on mathematical representations of complex physiological processes.

Drawing on more than three decades of research in cardiovascular mathematics, this talk highlights the fundamental contributions mathematicians have made to elucidating cardiovascular pathophysiology. Despite these advances, the translation of mathematical insight into routine clinical practice has historically been limited. The inherent complexity of cardiovascular systems, combined with cultural and structural barriers between disciplines, has constrained the clinical uptake of mathematical models and tools.

I will examine the reasons for this gap and argue that we are now at a pivotal moment. The unprecedented availability of high-quality clinical data, advances in data assimilation and uncertainty quantification, and the rapid development of artificial intelligence—when used in synergy rather than competition with mechanistic modeling and mathematical analysis—are transforming what is possible. Together, these developments open the door to clinically actionable models that can operate within realistic decision-making timelines.

A particular focus will be placed on ischemic heart disease and myocardial infarction, which currently account for approximately 400,000 deaths per year in the United States. I will argue that a substantial reduction in this number within the next five years is achievable, provided we succeed in translating existing mathematical knowledge into tools that clinicians can trust and use. This requires numerically efficient methods, grounded in strong theory, and designed explicitly for integration into clinical workflows.

The talk concludes with a forward-looking perspective: just as medical imaging revolutionized clinical practice a century ago, mathematics now has the opportunity—and responsibility—to drive the next transformation in cardiovascular care. The time is right for mathematicians to play a central role in shaping the future of clinical decision-making.

Student Talks

Student Talk Session I

Student Talk Session I	Friday, April 3	2:00–3:30 PM

Mitchell Scott Emory University

What Makes a Good Data Science Preconditioner?

Stochastic Gradient Descent (SGD) often slows in the late stage of training due to anisotropic curvature and gradient noise. We analyze preconditioned SGD in the geometry induced by a symmetric positive definite matrix $\mathbf{M}$. Our bounds make explicit how both the convergence rate and the stochastic noise floor depend on $\mathbf{M}$. For nonconvex objectives, we establish a basin-stability guarantee in a local $\mathbf{M}$-metric neighborhood around a minimizer set: under local smoothness and a local PL condition, we give an explicit lower bound on the probability that the iterates remain in the basin up to a time horizon. This perspective is particularly relevant in Scientific Machine Learning (SciML), where reaching small training losses under stochastic updates is closely tied to physical fidelity, numerical stability, and constraint satisfaction. Our framework covers both diagonal/adaptive and curvature-aware preconditioners and yields a practical criterion: choose $\mathbf{M}$ to improve local conditioning while attenuating noise in the $\mathbf{M}^{-1}$-norm. Experiments on a quadratic diagnostic and three SciML benchmarks support the predicted rate–floor behavior.

Jonathan Valyou Florida State University

A Conjugate-Gradient Formulation of the EnKF Algorithm

Data assimilation (DA) algorithms combine numerical forecast models with observation data to make accurate state estimations of a dynamical system in time. The Kalman Filter (KF) and Ensemble Kalman Filter (EnKF) are members of a widely popularized class of DA algorithms that uses a Bayesian approach to advance the state mean and covariance in time assuming linear Gaussian uncertainty conditions. These algorithms have a significant computational cost, in particular for highly dimensional dynamical systems. Algorithms have been established to reduce the computational complexity, yet there is typically a significant tradeoff with state estimation accuracy. In this work, we present a novel parallelizable reformulation of EnKF called the Conjugate Gradient Ensemble Kalman Filter (CGD-EnKF). We discuss analysis of an upper error bound for the CGD iterations, demonstrate CGD-EnKF's convergence properties, and showcase the algorithm's favorable computational efficiency–accuracy trade off in comparison with standardly utilized EnKF algorithms. Specifically, we will demonstrate that under certain conditions, the computational complexity of CGD-EnKF is of similar order as other standard EnKF algorithms while producing more accurate state estimations through numerical examples including those with high dimensionality in observation space.

Atticus Rex Georgia Tech

Multifidelity-Augmented Gaussian Process Inputs for Surrogate Modeling from Scarce Data

Supervised machine learning describes the practice of fitting a parameterized model to labeled input-output data. Supervised machine learning methods have demonstrated promise in learning efficient surrogate models that can (partially) replace expensive high-fidelity models, making many-query analyses, such as optimization, uncertainty quantification, and statistical inference, tractable. However, when training data must be obtained through the evaluation of an expensive model or experiment, the amount of training data that can be obtained is often limited, which can make learned surrogate models unreliable. However, in many engineering and scientific settings, cheaper low-fidelity models may be available, for example arising from simplified physics modeling or coarse grids. These models may be used to generate additional low-fidelity training data. The goal of multifidelity machine learning is to use both high- and low-fidelity training data to learn a surrogate model which is cheaper to evaluate than the high-fidelity model, but more accurate than any available low-fidelity model. This work proposes a new multifidelity training approach for Gaussian process regression which uses low-fidelity data to define additional features that augment the input space of the learned model. The approach unites desirable properties from two separate classes of existing multifidelity GPR approaches, cokriging and autoregressive estimators. Numerical experiments on several test problems demonstrate both increased predictive accuracy and reduced computational cost relative to the state of the art.

Francesco Brarda Emory University

Preconditioning Via Spectral Density Driven Graph Neural Networks

In this talk, we propose a learning-based framework for constructing sparse approximate inverse preconditioners for parameterized families of sparse linear systems arising from PDE discretizations. Our approach addresses two limitations of existing learned preconditioners: the mismatch between node-centric graph neural network architectures and entrywise preconditioner prediction, and the mismatch between Krylov convergence, which is governed by the spectrum of the preconditioned operator, and training objectives based only on algebraic surrogates such as Frobenius-norm losses. To address the first issue, we introduce a Line Graph with Virtual Node (LG-VN) architecture that promotes matrix nonzeros from edge attributes on the adjacency graph to node features on the line graph, enabling node-level regression aligned with row- and column-wise algebraic dependencies. To address the second, for symmetric positive definite systems we develop a differentiable spectral objective based on the Kernel Polynomial Method (KPM), which approximates the density of states of the symmetrically preconditioned operator using Chebyshev expansions and stochastic trace estimation. We combine this objective with a two-stage training strategy consisting of Frobenius-norm warm-up followed by spectral fine-tuning. Numerical experiments on diffusion-reaction problems show improved robustness and reduced Krylov iteration counts relative to classical FSAI and Frobenius-only training. We also report results on Helmholtz and convection-diffusion problems to examine the architecture beyond the SPD setting.

Haoyang Qian Florida State University

When Complexity Simplifies: The Paradox of More Random, More Tractable

Heterogeneity is intrinsic to real networked systems but fundamentally challenges the classical Master Stability Function (MSF), which relies on identical node dynamics and modal decoupling. When local dynamics vary across nodes, instabilities may emerge without a clear identification of which node is responsible. We introduce a localization-based generalized MSF that restores interpretability in heterogeneous networks. The approach exploits the tendency of Laplacian eigenvectors in large random networks to localize on small subsets of nodes. Each non-uniform mode can then be approximated by a reduced Jacobian governed by the local dynamics of the node on which the eigenvector is concentrated, yielding a heterogeneous MSF that predicts instability onsets while preserving a modal viewpoint. The uniform mode, which is not localized, is treated conservatively by evaluating node-wise Jacobians across the network. Crucially, the framework enables node-level attribution: when a small number of nodes are modified, the heterogeneous MSF identifies which specific node drives the instability. For illustration, we demonstrate this in a heterogeneous Brusselator network, where three perturbed nodes coexist with a large homogeneous background. Although all perturbed nodes contribute to spectral splitting, only one is correctly identified as responsible for the instability, in agreement with direct numerical simulations.

Student Talk Session II

Student Talk Session II	Friday, April 3	4:45–6:15 PM

Pengjun Wang Auburn University

Error Estimates of a Training-Free Diffusion Model for High-Dimensional Sampling

Score-based diffusion models are a powerful class of generative models, but their practical use often depends on training neural networks to approximate the score function. Training-free diffusion models provide an attractive alternative by exploiting analytically tractable score functions, and have recently enabled supervised learning of efficient end-to-end generative samplers. Despite their empirical success, the training-free diffusion models lack rigorous and numerically verifiable error estimates. In this work, we develop a comprehensive error analysis for a class of training-free diffusion models used to generate labeled data for supervised learning of generative samplers. By exploiting the availability of the exact score function for Gaussian mixture models, our analysis avoids propagating score-function approximation errors through the reverse-time diffusion process and recovers classical convergence rates for ODE discretization schemes, such as first-order convergence for the Euler method. Moreover, the resulting error bounds exhibit favorable dimension dependence, scaling as $\mathcal{O}(d)$ in the ℓ_2 norm and $\mathcal{O}(\log d)$ in the ℓ_∞ norm. Importantly, the proposed error estimates are fully numerically verifiable with respect to both time-step size and dimensionality, thereby bridging the gap between theoretical analysis and observed numerical behavior.

Caleb Fikes Emory University

PTYDIFF: Score-Induced Implicit Phase Choice Sampling for Nonlinear Ptychographic Imaging Problems

Many inverse problems exhibit intrinsic non-identifiabilities arising from symmetry groups of the forward model. In ptychography, a complex-valued object is reconstructed from intensity-only diffraction measurements corrupted by Poisson noise; because phase is unobserved, the posterior is invariant under global phase rotations. This gauge symmetry renders the model non-identifiable and induces a flat direction in the score field, which destabilizes diffusion-based posterior sampling in ambient coordinates. We introduce PtyDiff, a symmetry-aware diffusion framework that produces coherent posterior samples without external gauge fixing during inference. Our key contribution is Implicit Phase Choice (IPC), a training-time gauge-distance regularizer that softly selects a canonical representative within each symmetry orbit. We show that symmetry-induced posterior degeneracy in score-based generative models can be resolved by regularizing the score to select a smooth transversal of the symmetry orbits, thereby learning a symmetry-broken representation of the posterior in Euclidean space. Under mild regularity conditions, we prove that the resulting diffusion dynamics contract exponentially toward this learned section up to a noise-dependent residual, yielding globally phase-consistent samples. Empirically, PtyDiff achieves accurate reconstruction at as low as 20% probe overlap where classical rPIE fails, reducing NRMSE by up to 6x on in-distribution and 5x on out-of-distribution datasets, while recovering phase-consistent estimates directly without post-hoc alignment. The method generalizes to unseen object types and provides sampling-based uncertainty summaries, demonstrating robustness to mild distribution shift.

Jonathan Engle Florida State University

Elections and the Spatial Public Goods Game

The public goods game is a foundational framework for studying cooperation across various disciplines. Here, we model an electorate voting on the funding of a public good in a two-party system using evolutionary game theory. Voters adopt one of four strategies: consensus-makers, gridlockers, party 1 zealots, and party 2 zealots, which they may change via imitation. The public good benefits both local individuals and those in neighbouring regions due to spillover effects. A system of diffusion-advection-replicator equations governs the spatial movement of individuals and shifts in their voting strategies. Local social interactions drive strategy evolution, while migration occurs toward areas of higher utility, which is a function of both social and economic factors. Our results reveal bistability and significant spatial variations. Locally, populations converge to a gridlocked state or a mix of consensus-makers and zealots. We find that public good spillovers generate a free-rider effect, where poorly funded regions become spatially tied to, and dependent upon, well-funded ones.

Munawar Ali Florida State University

Noise Estimation of SDE from a Single Data Trajectory

In this talk, we propose a data-driven framework for model discovery of stochastic differential equations (SDEs) from a single trajectory, without requiring the ergodicity or stationary assumption on the underlying continuous process. By combining (stochastic) Taylor expansions with Girsanov transformations, and using the drift function's initial value as input, we construct drift estimators while simultaneously recovering the model noise. This allows us to recover the underlying $\mathbb{P}$ Brownian motion increments. Building on these estimators, we introduce the first *stochastic Sparse Identification of Stochastic Differential Equation (SSISDE)* algorithm, capable of identifying the governing SDE dynamics from a single observed trajectory without requiring ergodicity or stationarity. To validate the proposed approach, we conduct numerical experiments with both linear and quadratic drift–diffusion functions. Among these, the Black–Scholes SDE is included as a representative case of a system that does not satisfy ergodicity or stationarity.

Farhana Taiyebah Florida State University

Accelerating Posterior Inference from Pulsar Light Curves via Learned Latent Repre-

sentations and Local Simulator-Guided Optimization

Posterior inference from pulsar observations in the form of light curves is commonly performed using Markov chain Monte Carlo methods, which are accurate but computationally expensive. We introduce a framework that accelerates posterior inference while maintaining accuracy by combining learned latent representations with local simulator-guided optimization. A masked U-Net is first pretrained to reconstruct complete light curves from partial observations and to produce informative latent embeddings. Given a query light curve, we identify similar simulated light curves from the simulation bank by measuring similarity in the learned embedding space produced by pretrained U-Net encoder, yielding an initial empirical approximation to the posterior over parameters. This initialization is then refined using a local optimization procedure using hill-climbing updates, guided by a forward simulator, progressively shifting the empirical posterior toward higher-likelihood parameter regions. Experiments on the observed light curve of PSR J0030+0451, captured by NASA's Neutron Star Interior Composition Explorer (NICER), show that our method closely matches posterior estimates obtained using traditional MCMC methods while achieving 120 times reduction in inference time (from 24 hours to 12 minutes), demonstrating the effectiveness of learned representations and simulator-guided optimization for accelerated posterior inference.

Student Talk Session III

Student Talk Session III	Saturday, April 4	9:45–11:15 AM

Jingqiao Tang Florida State University

Data Assimilation Framework for Uncertainty Reduction in Learning Data-Driven Dynamical Systems

We introduce a score-filter-enhanced data assimilation framework designed to reduce predictive uncertainty in machine learning (ML) models for data-driven dynamical system forecasting. Machine learning serves as an efficient numerical model for predicting dynamical systems. However, even with sufficient data, model uncertainty remains and accumulates over time, causing the long-term performance of ML models to deteriorate. To overcome this difficulty, we integrate data assimilation techniques into the training process to iteratively refine the model predictions by incorporating observational information. Specifically, we apply the Ensemble Score Filter (EnSF), a generative AI-based training-free diffusion model approach, for solving the data assimilation problem in high-dimensional nonlinear complex systems. This leads to a hybrid data assimilation–training framework that combines ML with EnSF to improve long-term predictive performance. We Shall demonstrate that EnSF-enhanced ML can effectively reduce predictive uncertainty in ML-based Lorenz–96 system prediction.

Ruoyu Hu Florida State University

Introduction of Power Grid with a Numerical Example on Feedback Control Problem with Data Assimilation

Modern power grids are large-scale, chaotic dynamical networks. The operation is challenged by lots of factors such as extreme events and increasing demand of electricity. This talk introduces the power grid setting and the Grid2op/L2rpn simulation framework, then presents a numerical example for studying data assimilation in improving the control of the simulated power grid.

Yi Liu Auburn University

Multi-level Machine Learning Framework for Inverse Scattering Problems

Inverse medium and inverse source scattering problems with finite measurements are ill-posed and highly sensitive to measurement noise. In this talk, we present a multi-level deep learning framework for reconstruction from multi-frequency boundary measurements. The method combines multi-frequency iterative regularization with multi-grade deep learning, in which the reconstruction is refined progressively. We further analyze the learning dynamics and generalization behavior in the Neural Tangent Kernel (NTK) regime. Numerical experiments, including comparisons with Landweber-type iterative baselines, show that the proposed approach is stable with respect to measurement noise and leads to improved reconstruction quality.

Tomoki Koike Georgia Tech

Sparse POD Mode Selection and Manifold Dimensionality Reduction with Neural Networks

High-performance simulations can resolve complex physical systems in high fidelity, but tasks like control and inverse problems remain computationally prohibitive, motivating the need for fast, accurate reduced-order models. A standard tool, Proper Orthogonal Decomposition (POD), builds these models by projecting dynamics onto a low-dimensional linear subspace. Yet for advection-dominated and turbulent flows, POD requires many modes to achieve accurate reconstruction, and its energy-based mode selection can discard physically important small-scale features. Recent work on nonlinear manifold methods addresses some of these shortcomings, but existing approaches either fix the nonlinear mapping form ahead of time (limiting flexibility) or rely on neural networks with energy-based selection, sacrificing interpretability. In this talk, we present SparseModesNet, a dimensionality reduction framework that marries the interpretability of POD with the expressivity of neural network decoding. The key ingredient is LassoNet, a method which enforces hierarchical sparsity via residual connections, enabling simultaneous data-driven mode selection and nonlinear reconstruction. We demonstrate SparseModesNet on benchmark advection-dominated and chaotic flow problems, where it matches or exceeds state-of-the-art reconstruction accuracy. On turbulent channel flow with friction Reynolds number of 5200, it cuts reconstruction error by 51–78% relative to existing state-of-the-art methods, while retaining interpretable, physically meaningful mode selection.

Seungmin Lee Florida State University

An optimization-based bound-preserving limiter for two-phase systems

Multiphase flows in porous media appear in a large number of applications in engineering and science. One of the issues arising in modeling and solving a two-phase flow system is the violation of the maximum principles in the saturation transport equation. To ensure accurate and physically consistent numerical approximations, bound-preserving methods have been suggested and well studied in the literature. In this work, we introduce an optimization-based bound-preserving limiter for solving a two-phase flow system. This algorithm enforces bounds on the cell averages of DG solutions so that the point values can easily be enforced to be bounded by a simple scaling limiter on the DG polynomial in each cell. The optimization can be efficiently solved by a first-order splitting method with nearly optimal parameters, which has O(N) computational complexity. Numerical tests are shown for some examples to demonstrate the performance of the method.

Student Talk Session IV

Student Talk Session IV	Saturday, April 4	1:45–3:15 PM

Ferhat Karabatman Florida State University

Geometric Perspective on Concentration Phenomena in Frame Theory

Frames are fundamental structures in many areas, and tight frames are particularly valued for their stability and robustness properties. In this work, we establish concentration phenomena for Parseval frames, i.e. tight frames with frame bound 1, under isotropic distributions supported on the sphere and the Euclidean ball, showing that epsilon-nearly Parseval frames are prevalent in these probabilistic models. We further introduce a distinguished subclass of Parseval frames and prove that they are both robust under the Bernoulli-type erasure model and prevalent within the space of Parseval frames. As an application of our results, we obtain a high-probability upper bound of order $O(\epsilon^2 d)$, which achieves the optimal order for the Paulsen problem.

Xu (Melissa) Wang Emory University

A Primal-Dual Price-Optimization Method for Computing Equilibrium Prices in Mean-Field Games Models

In this work, we consider a mean-field model of electricity market where households make decentralized buy-sell decisions based on the electricity price. The equilibrium price corresponds to a market-clearing condition in which electricity demand and supply balance perfectly, reducing strain on the grid. The main challenge in this model is finding the equilibrium price, which is often achieved by solving a coupled system of nonlinear PDEs. Therefore, we propose a new, simple, modular, yet efficient primal-dual optimization algorithm based on the variational formulation of the price.

Mujtaba Ali Florida State University

On the Hausdorff Stability of Barcodes over Posets

The Isometry Theorem for single-parameter persistence modules is a fundamental result in TDA, stating that the bottleneck distance and interleaving distance are equal between the persistence modules. Extending this result to multiparameter persistence modules is significantly more challenging. In higher dimensions, persistence modules are generally not interval decomposable, so the classical barcode representation does not always exist. To address this, researchers often restrict attention to classes of modules that admit interval decompositions, allowing a notion of bottleneck distance to be defined. However, even within this restricted setting, the relationship between bottleneck and interleaving distances behaves very differently. Lesnick and Botnan conjectured that a Lipschitz bound holds when modules are convex interval decomposable, and Bjerkevik later proved this conjecture for the special case of rectangle-decomposable persistence modules. In our work, we consider the Hausdorff distance on barcodes and show that, for convex interval decomposable multiparameter persistence modules, the barcode map is 2-Lipschitz with respect to this distance.

Jiaqi Yang Emory University

Multi-patient Computational Analysis of Type B Aortic Dissections

Type-B Aortic Dissection (TBAD) is a life-threatening cardiovascular pathology in which a primary intimal tear in the descending aorta creates a true lumen and false lumen (FL). Predicting FL progression remains a critical clinical challenge, yet existing studies relying on morphological biomarkers from computed tomography angiography and hemodynamic parameters from 4D flow MRI and computational fluid dynamics (CFD) simulations have yielded inconsistent results, largely due to small cohort sizes and complex patient-specific geometry. Key hemodynamic indices including time-averaged wall shear stress (TAWSS), oscillatory shear index (OSI), relative residence time (RRT), and the topological shear variational index (TSVI) are investigated alongside morphological features for their association with FL growth. By integrating both morphological and hemodynamic analyses, this study aims to identify robust predictive biomarkers for FL expansion in TBAD, contributing toward improved risk stratification and more informed clinical decision-making.

Haoran Yan Georgia Tech

Understanding Denoising Autoencoders through the Manifold Hypothesis: A Geometric Perspective

Denoising autoencoders are a standard architecture to recover clean signals from noisy high-dimensional data, but their performance is often better explained by low-dimensional geometry. In this talk, I will present a geometric perspective based on the manifold hypothesis, where clean data are assumed to lie near a smooth low-dimensional manifold, and the observations are generated by adding isotropic Gaussian noise to clean samples. We will see that the denoising error is expected to scale mainly with the intrinsic dimension d, rather than with the ambient dimension D. I will outline the theoretical intuition behind this phenomenon and present numerical experiments on synthetic manifolds and image datasets, including MNIST and Yale faces. Empirically, we observe stable log-log scaling of test MSE with sample size, with exponents consistent with low-dimensional structure and intrinsic-dimension estimates computed directly from the data.

Student Talk Session V

Student Talk Session V	Saturday, April 4	3:45–5:15 PM

Yi-Yung Yang
Florida State University

Bound-preserving and Entropy Stable Enriched Galerkin Methods for Nonlinear Hyperbolic Equations

In this talk, we present monolithic limiting techniques for enforcing nonlinear stability constraints in enriched Galerkin (EG) discretizations of nonlinear scalar hyperbolic equations. Local mass conservation and control over cell averages are achieved by enriching continuous (multi-)linear finite element spaces with piecewise-constant functions, resulting in a spatial semi-discretization with a variational multiscale structure. While the method is inherently stable for linear advection, it is generally not bound-preserving. To enforce discrete maximum principles and entropy stability in the nonlinear case, we apply limiters adapted to the locally conservative EG formulation. The proposed algorithms build on recent advances in convex limiting and algebraic entropy fixes for finite element methods. Numerical experiments for two-dimensional nonlinear hyperbolic problems demonstrate that the limiters effectively enforce the imposed constraints while preserving optimal accuracy for smooth solutions.

Kha Doan
Auburn University

Dynamically Regularized Lagrange Multiplier Method for the Cahn-Hilliard–Navier–Stokes System

This work is concerned with efficient and accurate numerical schemes for the Cahn-Hilliard-Navier-Stokes phase field model of binary immiscible fluids. By introducing two Lagrange multipliers for each of the Cahn-Hilliard and Navier-Stokes parts, we reformulate the original model problem into an equivalent system that incorporates the energy evolution process. Such nonlinear, coupled system is then discretized in time using backward differentiation formulas, in which all nonlinear terms are treated explicitly and no extra stabilization term is imposed. The resulting Dynamically Regularized Lagrange Multiplier (DRLM) schemes are fully decoupled, mass conserving, and unconditionally energy stable with respect to the original variables. Numerical experiments demonstrate the accuracy and robustness of the proposed DRLM schemes.

Leonardo Molinari
Emory University

Computational Framework for Cardiac Radiofrequency Ablation: Multiphysics Modeling and Domain Decomposition

Radiofrequency ablation (RFA) is a cornerstone treatment for cardiac arrhythmias, yet its long-term efficacy is often limited by a lack of understanding regarding the biophysical mechanisms of lesion formation. Traditional models rely on thermal thresholds that fail to distinguish between reversible tissue stunning and the permanent functional block required for clinical success.

This talk introduces a high-fidelity computational framework that integrates cardiac electrophysiology (EP) directly into RFA assessment. By shifting the paradigm toward functional conduction block, this work establishes a new foundation for predictive ablation modeling.

The framework, implemented within the MFEM finite element library, comprises two pillars:

(I) A Multiphysics RFA Solver that couples electrostatics, bioheat transfer, and fluid dynamics using efficient domain decomposition.

(II) An Integrated EP Solver that utilizes an automated ODE code-generation pipeline to handle complex membrane kinetics.

The primary novelty lies in the dynamic coupling between these solvers, where thermal feedback modulates tissue excitability in real-time. By leveraging the DOE-supported software, this research demonstrates how advanced numerical methods can optimize ablation strategies in silico, paving the way for personalized treatment planning and improved patient outcomes.

Akshita Sahni Georgia Tech

A Reduced Order Model for Rapid 3D Fluid-Structure Interaction Simulation of Aortic Valve Flows

Computational fluid dynamics is becoming prevalent for pre-surgical planning of complex cardiac surgeries. Fluid-structure interaction simulations such as flow through aortic valves can take one day to simulate transvalvular flow dynamics for just one cardiac cycle on a supercomputing cluster. The high computational cost makes it less feasible to explore patient-specific parameter variations through multiple simulations, within the short clinical timelines of pre-surgical planning. Model order reduction is a powerful technique that can reconstruct key hemodynamic features while reducing the computational cost by orders of magnitude. We develop a 3D reduced order model (ROM) of aortic transvalvular flow based on fully coupled FSI simulations. Our ROM prediction + velocity field reconstruction of peak systolic time step required 0.0415s on a 48GB RAM laptop, achieving $\sim 10{,}000\times$ speedup. The ROM recovered dominant transient flow structures over peak systole. Because the training data are fully coupled FSI simulations, the ROM predictions reflect coupled valve dynamics. With this project, our goal is to make pre-surgical planning of complex heart valve surgeries feasible for short clinical turnaround times.

James Orgeron Florida State University

Habitat Fragmentation Promotes Spatial Scale Separation Under Resource Competition

Habitat fragmentation, often driven by human activities, alters ecological landscapes by disrupting connectivity and reshaping species interactions. In such fragmented environments, habitats can be modeled as networks, where individuals disperse across interconnected patches. We consider an intraspecific competition model, where individuals compete for space while dispersing according to a nonlinear random walk, capturing the heterogeneity of the network. The interplay between asymmetric competition, dispersal dynamics, and spatial heterogeneity leads to nonuniform species distribution: individuals with stronger competitive traits accumulate in central (hub) habitat patches, while those with weaker traits are displaced toward the periphery. We provide analytical insights into this mechanism, supported by numerical simulations, demonstrating how competition and spatial structure jointly influence species segregation. In the large-network limit, this effect becomes extreme, with dominant individuals disappearing from peripheral patches and subordinate ones from central regions, establishing spatial segregation. This pattern may create favorable conditions for speciation, as physical separation can reinforce divergence within the population over time.

Posters

Kapil Chawla Florida State University

DG-FEONet: Learning Discontinuous PDE Solutions with Neural Networks

We introduce Discontinuous Galerkin Finite Element Operator Network (DG-FEONet), a data-free operator learning framework for solving parametric partial differential equations (PDEs) with discontinuous coefficients and non-smooth solutions. Traditional neural operator methods such as DeepONet and Fourier Neural Operator typically require large paired datasets and often struggle to accurately capture sharp solution features arising from discontinuities. Our approach combines the discontinuous Galerkin (DG) finite element method with neural networks by minimizing the residual of a DG-based weak formulation using the Symmetric Interior Penalty Galerkin (SIPG) scheme. The neural network predicts element-wise solution coefficients, allowing the model to be trained without precomputed input–output solution pairs. We provide theoretical justification through convergence analysis and demonstrate the effectiveness of the method on several one- and two-dimensional PDE problems. The results show that DG-FEONet accurately captures discontinuities, generalizes well across parameter variations, and provides reliable convergence behavior. This work highlights the potential of integrating local numerical discretization methods with machine learning to achieve robust operator learning for challenging PDE problems.

Srijon Sarkar Emory University

Kronecker Approximations of Covariance Matrices for Solving Inverse Problems

Inverse problems arise in many applications where the task is to recover unknown parameters or data from indirect or noisy observations. In contrast to forward problems, where outputs are directly computed from known inputs, inverse problems are designed to infer the underlying variables for the observed effect. In many settings, large prior covariance matrices pose computational bottlenecks for solving inverse problems. To address this, we exploit Kronecker product approximations and demonstrate how these approximations can be used to solve inverse problems.

Ngoc Hien Tran Auburn University

Reduced Fracture Model for Reactive Transport

Reactive transport refers to the study of modelling underground flow, which plays a vital role in computational geochemistry and has various applications. The mathematical model couples the transport process, described by advection-diffusion PDEs, with the chemical model, which is either ODEs or algebraic equations, depending on whether the reactions involved are of kinetic or equilibrium type. The real-world subsurface, in many cases, has fractures and faults that can speed up or slow down the geochemical process, resulting in more complex flow modelling. We study the reduced fracture model, in which the fractures are treated as interfaces between different subdomains. The well-posedness analysis of the model is carried out, and preliminary numerical results for multiple test cases are shown.

Yahong Yang Georgia Tech

Learning in Group-Invariant Radon Spaces from Finite Samples

In this talk, we first show that, when learning a Radon-TV function with a group-invariant structure from finite samples, any global minimum corresponds to a shallow neural network that also exhibits the same group-invariant structure, following the work of Journal of Machine Learning Research (2021, 22(43): 1–40). We then study the generalization behavior of learning in group-invariant Radon-TV spaces. In particular, we provide an upper bound on the generalization error and a corresponding upper bound on the minimax rate, both of which highlight the benefits of group invariance in certain regimes. The upper bound suggests that group invariance can yield significant gains in low-dimensional settings; however, in higher dimensions, since Radon-TV spaces already exhibit low-dimensional structure, the additional advantage from group invariance is reduced in this bound. On the other hand, for the lower bound (i.e., the minimax rate), we show that incorporating group invariance leads to lower metric entropy of the Radon-TV space, indicating that group-invariant structure indeed reduces the intrinsic learning complexity.

Vivian Zhang Georgia Tech

Multifidelity Operator Inference: Reduced Order Modeling of PDEs from Scarce Data

Large-scale engineering simulations, governed by partial differential equations (PDEs), are expensive to compute. One solution is to use operator inference (OpInf), a scientific machine learning and model reduction approach, to approximate reduced-state operators for PDEs. OpInf utilizes high-dimensional dynamical data to non-intrusively construct a low-dimensional model such that the original physics are preserved. While OpInf simulates high-dimensional system dynamics with great accuracy, the approach requires access to high-fidelity data, which in practice is scarce, resulting in poorly trained models with high variance. We propose a new multifidelity approach to OpInf that exploits the context where data and models of varying costs and fidelities are available. Utilizing an approximate control variate framework, we define multifidelity Monte Carlo estimators for the OpInf reduced-order operators and provide numerical results from a linear-quadratic PDE example to show a reduction in variance in the multifidelity approach compared to the single, high-fidelity approach.

Shikhar Shah Emory University

Block Jacobi Preconditioning on Analog Hardware

Analog in-memory computing leverages non-volatile resistive memory to perform linear algebra operations via Ohm's and Kirchhoff's laws. These operations are inherently approximate and stochastic; however, key kernels like matrix-vector products execute in time independent of the matrix size, making analog hardware attractive for large, sparse linear systems. We present convergence results for a block Jacobi preconditioner on analog hardware, comparing construction methods including the sparse approximate inverse and the Monte Carlo approximate inverse, with evaluation on the Laplace and convection-diffusion-reaction equations. We find that optimal preconditioner parameters differ substantially from their digital counterparts, and further extend our approach to a hierarchical two-level preconditioner scalable to very large systems.

Anand Natarajan Sriram Georgia Tech

Multifidelity Gaussian Processes for Vortex Dominated Flows

Studies on vortex-dominated flows typically require computationally expensive simulations in order to resolve all scales in the flow-field and accurately estimate the dynamics of the system being studied. This work presents the use of Gaussian Processes to combine few computationally expensive and accurate (high-fidelity) simulations and multiple computationally cheap but potentially inaccurate (low-fidelity) simulations for a pitching airfoil system. The multi-fidelity framework uses nonlinear autoregressive Gaussian Processes to combine these model fidelities and provide significant reductions in evaluations of the high-fidelity model while retaining the accuracy of the system as well as provide estimates of the model uncertainty.

Benjamin Burns Georgia Tech

Infinite-dimensional Stein variational inference with derivative-informed neural operators

Bayesian inversion of parametric PDEs is a powerful tool for quantifying uncertainty in scientific models, but its application is often limited by the high cost of repeated PDE solves in high-dimensional settings. We present a surrogate-assisted projected Stein variational gradient descent (pSVGD) framework that addresses these challenges. Our approach integrates data-driven model reduction with a derivative-informed neural operator that learns the PDE solution operator in reduced space, amortizing the cost of PDE solves across inference iterations. By combining parameter and solution projections with accurate, derivative-informed surrogate training, we obtain efficient and reliable score function evaluations for particle updates. Theoretical error bounds provide control of surrogate accuracy, while numerical results demonstrate substantial computational savings without loss of inference fidelity. This framework opens the door to fast and scalable Bayesian inference for complex PDE models with high-dimensional parameters.

Meixi Li Georgia Tech

Detecting repeating patterns in signals

We propose a general framework to find nonperiodic repeating patterns in any given signal. A response vector is built at each point using the self-similarity of the given signal. The regularized k-means clustering is applied to group the response vectors, through which dimensionality reduction is achieved. Based on the cluster labels at each point, we find all repeating patterns by extracting segments with the same sets of labels.

Tianshi Xu Emory University

Preconditioned Truncated Single-Sample Estimators for Scalable Stochastic Optimization

Many large-scale stochastic optimization problems require repeated solutions of linear systems and evaluations of log-determinants. While exact computation is often prohibitively expensive, classical truncated iterative methods introduce bias, and existing unbiased Krylov estimators may exhibit high variance and numerical instability. Preconditioned Truncated Single-Sample (PTSS) estimators address these challenges by integrating preconditioning with truncated Lanczos iterations. This approach produces stable, low-variance estimators for linear system solutions, log-determinants, and their derivatives. Both theoretical analysis and numerical experiments demonstrate that PTSS improves robustness, reduces variance, and increases computational efficiency in large-scale stochastic optimization.

Toma Debnath Florida State University

On the emergence of antiphase synchronization

Pattern formation in network-coupled dynamical systems plays an important role in many physical, chemical, and biological processes. In this work, we investigate the emergence of spatial patterns in reaction–diffusion dynamics on networks, with a particular focus on quasi-ring topologies. The dynamics are studied using a two-species reaction–diffusion model on a network, where diffusion is represented through the graph Laplacian. The stability of the homogeneous steady state is analyzed using the Master Stability Function framework. By decomposing perturbations into Laplacian eigenmodes, the stability of each spatial mode can be determined from the associated Jacobian spectrum. To understand the nonlinear development of instabilities, weakly nonlinear theory is employed, leading to amplitude equations that describe the growth and saturation of patterns near the instability threshold. Numerical simulations of the Brusselator model are performed on both regular ring networks and perturbed quasi-ring networks. While the symmetric ring supports spatially extended patterns associated with delocalized Fourier modes, small perturbations to the coupling weights break this symmetry and produce localized Laplacian eigenvectors. In the quasi-ring topology, the eigenvectors that become localized are those associated with the short-wave instability, and their interaction with the delocalized modes responsible for the long-wave instability plays a key role in shaping the resulting patterns. This eigenvector localization strongly influences the resulting dynamics and can confine activity to specific regions of the network. As a consequence, the quasi-ring network exhibits a variety of dynamical regimes, including synchronized oscillations, oscillation death, traveling waves, and mixed spatial states in which stationary and oscillatory regions coexist. These results demonstrate how small structural perturbations can significantly alter the spectral

properties of a network and lead to complex pattern formation, including chimera-like behavior.

Zhongjie Shi Georgia Tech

Learning theory of sequence-to-sequence mapping between manifolds with Transformers

Abstract not provided.

Callie Reid Florida State University

Dynamics Shape the Perception of Network Structure

Dynamical processes can fundamentally reshape how network structure is perceived, with functional organization emerging even when topological modularity is weak or absent. We investigate this phenomenon in nonlinear random walks—mean-field models of diffusion with volume exclusion that are nonlinear on heterogeneous networks and linear otherwise. For connected undirected graphs, we analyze the rate of convergence to equilibrium for n interacting walkers and introduce a localization-based reduction exploiting the spontaneous concentration of Laplacian eigenvectors on specific nodes in large random graphs. We show analytically, and validate numerically, that in scale-free networks the Jacobian spectrum of the nonlinear random walk undergoes a characteristic spectral deformation: eigenvalues associated with peripheral, low-degree nodes are compressed toward the origin, while those linked to hubs are systematically displaced outward. This nonlinear spectral reorganization reveals how heterogeneity and crowding jointly redefine the network's effective dynamical structure, offering a principled mechanism for dynamic modularity in complex systems. Building on this perspective, we introduce weighted-node paths in which nodes are weighted according to their crowding capacity, allowing dynamical proximity to be quantified through path-based measures that naturally extend classical closeness and betweenness centrality to account for congestion effects induced by the dynamics.

Shiyi Lyu Florida State University

Improvement of Reducing Dynamical Systems on Networks

Low-dimensional reductions provide an important tool for understanding spreading dynamics on complex networks. We study nonlinear SIS-type dynamics on undirected scale-free networks, where strong degree heterogeneity can drive sharp transitions. We focus on a reversible spreading model with nonlinear incidence on a weighted adjacency network and use forward and backward parameter sweeps to numerically resolve bistability and identify the transition points. Building on the Gao–Barzel–Barabási (GBB) mean-field reduction, we use the GBB formulation as a baseline and then refine the reduction to better capture the heterogeneous core structure typical of scale-free graphs. This refinement reduces systematic offsets in predicted thresholds and more accurately reproduces the branch structure observed near explosive transitions. The resulting reduced description enables improved prediction of critical points and transition behavior compared with the standard GBB reduction. By contrast, in Erdős–Rényi networks, where heterogeneity is much weaker, the standard GBB reduction already provides an accurate description and no refinement is required.

Zixiang Xiong Florida State University

Applying Linear Regression on Ensemble score filter

Recent advances in data assimilation (DA) have focused on developing more flexible approaches that can better accommodate nonlinearities in models and observations. However, it remains unclear how the performance of these advanced methods depends on the observation network characteristics. In this study, we present initial experiments with the surface quasi-geostrophic model, in which we compare a recently developed AI-based ensemble filter with the standard Local Ensemble Transform Kalman Filter (LETKF). Our results show that the analysis solutions respond differently to the number, spatial distribution, and nonlinear fraction of assimilated observations. We also find notable changes in the multiscale characteristics of the analysis errors. Given that standard DA techniques will be eventually replaced by more advanced methods, we hope this study sets the ground for future efforts to reassess the value of Earth observation systems in the context of newly emerging algorithms.

Benjamin Yellin Emory University

First-Order PINNs and LSFEM

Physics-Informed Neural Networks (PINNs) have proven to be very useful for solving partial differential equations, particularly as a meshless alternative to the Finite Element Method. However, they can be difficult to train due to the non-convex loss landscape. We draw a connection between the first-order formulation of PINN training and the Least-Squares Finite Element Method. We demonstrate how the particular choice of first-order formulation of the advection-diffusion equation impacts the performance of our method and relate this to the norm-equivalence of the corresponding least-squares functional.