MID-ATLANTIC NUMERICAL ANALYSIS DAY

A one-day conference on Numerical Analysis and Scientific Computing for graduate students and post-doctorates from the Mid-Atlantic region

Department of Mathematics

Friday, November 14, 2025 | Philadelphia, PA

Organizers: Benjamin Seibold and Daniel B. Szyld

Sponsored by the Department of Mathematics,
College of Science and Technology, The Graduate School, the Center for Computational Mathematics and Modeling, Temple University, and the Simons Foundation

Temple University

Schedule (Friday, November 14, 2025)

9:15am-9:50am	Registration and Breakfast (provided)
9:50am 10:00am	Opening Remarks
10:00am-11:00am	Presentations (Time-Stepping)
11:00am 11:20am	*Coffee Break*
11:20m-12:20pm	Presentations (Low-Rank Methods)
12:20pm 1:30pm	Lunch (provided)
1:30pm-2:30pm	Keynote lecture (Mayya Tokman)
2:30pm-3:00pm	*Coffee Break*
3:00pm-4:00pm	Presentations (Applications and CFD)
4:00pm-4:20pm	*Coffee Break*
4:20pm-5:00pm	Presentations (Machine Learning)
5:00pm-5:10pm	Closing Remarks
6:00pm-8:00pm	Group Dinner (attendance optional)

Speakers

Time-Stepping (10:00am-11:00am)

Sylvia Amihere, UNIVERSITY OF MARYLAND, BALTIMORE COUNTY

Embedded Implicit-Explicit Strong Stability Preserving Runge-Kutta Methods

Isaac Anthony Castro, UNIVERSITY OF DELAWARE

Low-Rank Fully Implicit Runge-Kutta Method for Linear PDEs

Yifan Hu, UNIVERSITY OF MARYLAND, BALTIMORE COUNTY

Adaptive-Step Time Integration for Gyrokinetic Equation with GENE-X

Low-Rank Methods (11:20am–12:20pm)

Jackie Lok, PRINCETON UNIVERSITY

Sketch-and-Project Solvers for Linear Systems with Low-Rank Structure

Mikhail Lepilov, RENSSELAER POLYTECHNIC INSTITUTE

Kernel Approximation Using the Proxy Point Method via Contour Integration

Shambhavi Suryanarayanan, PRINCETON UNIVERSITY

On Trimming Tensor-Structured Measurements and Efficient Low-Rank Tensor Recovery

Keynote Lecture (1:30pm-2:30pm)

Mayya Tokman
University of California, Merced

Title: Why, What, and How of Exponential Integration

Applications and CFD (3:00pm–4:00pm)

Ahmet Kaan Aydin, UNIVERSITY OF MARYLAND, BALTIMORE COUNTY

Low-Rank All-at-Once Stochastic Galerkin Solvers for Incompressible Flows

Arnab Roy, UNIVERSITY OF DELAWARE

Inverse Modeling of Tear-Film Thinning: From ODE Dynamics to Neural Networks

Alex Sheranko, UNIVERSITY OF MARYLAND, BALTIMORE COUNTY

A Kernel-Based SIR Model for Infectious Spread in Space and Time A kernel-based SIR model

Machine Learning (4:20pm–5:00pm)

Noah Amsel, NEW YORK UNIVERSITY

The Polar Express: Optimal Matrix Sign Methods and Their Application to the Muon Algorithm

Leonardo Ferreira Guilhoto, UNIVERSITY OF PENNSYLVANIA

Kolmogorov Superposition Meets Physics-Informed Neural Networks: A New Architecture for Scientific Machine Learning

Abstracts of Talks

Time-Stepping (10:00am-11:00am)

Sylvia Amihere, University of Maryland, Baltimore County

Embedded Implicit-Explicit Strong Stability Preserving Runge-Kutta Methods

Strong stability preserving Runge-Kutta methods are high-order time stepping methods that are widely used to solve time-dependent partial differential equations (PDEs). IMEX SSP Runge-Kutta methods allow efficient treatment of PDEs that are split into a stiff part (treated implicitly) and a non-stiff part (treated explicitly), while still ensuring preservation of stability properties, such as monotonicity or total variation diminishing behavior, under an appropriate time-step restriction. In this work, we present embedded IMEX SSP Runge-Kutta methods that use two schemes of different orders of accuracy sharing the same stages. The embedded IMEX SSP pairs allow for efficient error estimation and adaptive step size control while preserving strong stability. We discuss the construction of such embedded pairs, their SSP coefficients, and accuracy properties. Numerical tests on representative PDEs demonstrate the benefits of the embedded structure for adaptive algorithms. Benchmarks compare the performance of embedded IMEX SSP schemes with non-embedded counterparts and standard IMEX schemes to analyze the efficiency and robustness of these new methods.

Isaac Anthony Castro, University of Delaware

Low-rank fully-implicit Runge-Kutta method for linear PDEs

Fully implicit Runge-Kutta methods have high-order accuracy and excellent stability properties for stiff problems, but their computational cost can be prohibitive due to large coupled systems at each step. This talk introduces a low-rank approach that exploits the structure of linear PDE discretization to reduce the complexity of solving the coupled stage systems. We present a formulation that represents the stage vectors on a low-rank basis and applies iterative solvers with low-rank preconditioners to accelerate convergence. Numerical tests on linear addiction-diffusion problems show that the proposed method significantly reduces wall-clock time while preserving accuracy and stability of the original fully implicit scheme. We also discuss extensions to multi-dimensional problems and performance on benchmark flows, including comparisons with classical diagonally implicit Runge-Kutta (DIRK) schemes and semi-implicit approaches. Applications highlight the method's potential in large-scale simulations, with an outlook toward nonlinear problems via linearization strategies such as Newton-Krylov methods and a streamline-vorticity formulation of the Navier-Stokes equations.

Yifan Hu, University of Maryland, Baltimore County

Adaptive-Step Time Integration for Gyrokinetic Equation with GENE-X

We present progress on the implementation of adaptive-step time integration for the GENE-X gyrokinetic code. The gyrokinetic equation is a five-dimensional PDE that models plasma behavior in fusion devices and is computationally intensive due to multiscale dynamics and stiffness. Our approach integrates the SUNDIALS time integration library to provide reliable adaptive step-size control using embedded methods and error estimation. We describe the coupling of GENE-X with SUNDIALS, focusing on efficient Jacobian-vector products, preconditioning strategies tailored to the gyrokinetic operator, and handling of stiff collision terms. Numerical results for the TCV-X21 case study demonstrate improved efficiency and robustness compared to fixed-step integrators, with automatic time-step selection that adapts to transient behavior and stiff regimes. The implementation reduces the need for manual tuning of time steps and allows for larger, problem-tailored steps while meeting user-specified tolerances. The framework supports future extensions to implicit-explicit (IMEX) integrators and advanced error controllers that leverage invariants and physics-informed metrics. In this talk, we present an overview of GENE-X and SUNDIALS integration, discuss performance considerations on modern architectures, and show preliminary results from adaptive integration for the 5D gyrokinetic nuclear fusion case TCV-X21.

Low-Rank Methods (11:20am–12:20pm)

Jackie Lok, Princeton University

Sketch-and-project solvers for linear systems with low-rank structure

The sketch-and-project method is a unifying framework for many randomized iterative methods for solving linear systems. In this talk, we introduce a new family of sketch-and-project solvers that exploit low-rank structure inherent in the coefficient matrix or right-hand side. Our analysis shows that when the system admits an approximate low-rank representation, properly chosen sketches dramatically accelerate convergence while maintaining robustness to noise and ill-conditioning. We develop adaptive sketching strategies that estimate and track the effective numerical rank during iteration, enabling automatic adjustment of sketch size for efficiency. We also extend the framework to block systems and matrix equations with Kronecker structure. Experiments on synthetic and application-driven problems, including PDE-derived systems and machine learning kernels, demonstrate consistent speedups over classical randomized Kaczmarz and sketch-and-project variants. The methods are simple to implement, parallel-friendly, and compatible with preconditioning. We provide theoretical guarantees on convergence rates under realistic assumptions and discuss connections to Nystrom approximations, showing how sketch selection can be aligned with capturing the subspace that corresponds to a good low-rank Nystrom approximation.

Mikhail Lepilov, Rensselaer Polytechnic Institute

Kernel approximation using the proxy point method via contour integration

Large kernel matrices arise frequently in several areas of numerical analysis and data science, but forming and factorizing them is computationally expensive. We introduce a proxy point method based on contour integration to approximate kernel matrices with controlled accuracy. The approach represents far-field interactions using proxy sources placed on complex contours, enabling efficient low-rank approximations without sampling the full matrix. We provide an error analysis that quantifies the trade-off between the number and placement of proxy points and the resulting approximation quality. The method applies broadly to kernels with analytic structure, including Laplace and Gaussian kernels, and can be incorporated into hierarchical matrix formats. Numerical experiments show competitive performance against state-of-the-art kernel approximation schemes in terms of runtime and memory usage, while achieving the best available theoretical bounds on approximation accuracy.

Shambhavi Suryanarayanan, Princeton University

On trimming tensor-structured measurements and efficient low-rank tensor recovery

In this talk, we will consider the problem of low-rank tensor recovery from tensor-structured measurements. We propose a new trimming procedure on the measurement operator that improves the conditioning of the recovery problem and leads to provably efficient algorithms for the reconstruction of low-rank tensors from incomplete or corrupted observations. Our framework leverages tensor algebra and randomized sketching to reduce dimensionality while preserving the essential multilinear structure. We provide recovery guarantees and sample complexity bounds under standard incoherence-type conditions and analyze robustness to noise and outliers. Numerical experiments demonstrate that trimming substantially improves convergence rates and reconstruction quality compared to untrimmed baselines, especially in ill-conditioned regimes. We conclude with extensions to tensor completion and robust recovery settings, highlighting the implications for large-scale scientific data and CP-rank tensors. This is joint work with Elizaveta Rebrova.

Applications and CFD (3:00pm–4:00pm)

Ahmet Kaan Aydin, University of Maryland, Baltimore County

Low-rank all-at-once stochastic Galerkin solvers for incompressible flows

We study iterative solvers for unsteady incompressible Navier Stokes equations with stochastic input parameters using an all at once stochastic Galerkin formulation. The resulting linear systems are extremely large due to coupling across steps and stochastic modes. We develop low-rank Krylov subspace methods that exploit separability in space, time, and parameter dimensions, combined with tailored block preconditioners that preserve low-rank structure. The approach reduces both memory footprint and computational cost, enabling efficient uncertainty quantification for realistic flow problems. We present numerical experiments on canonical benchmark cases, including driven cavity and channel flows with random viscosity, showing significant speedups and accurate statistical estimates of quantities of interest. The methodology is compatible with hierarchical tensor formats and parallel computing and provides a foundation for robust all at once solvers in high dimensional stochastic PDEs. The effectiveness of the method is illustrated by numerical experiments.

Arnab Roy, University of Delaware

Inverse modeling of tear-film thinning: From ODE dynamics to neural networks

The tear film is a thin fluid layer that coats and protects the ocular surface. Its thinning dynamics are governed by complex coupled processes including evaporation, osmotic flow, Marangoni effects, and lipid layer behavior. We address the inverse problem of estimating key physical parameters driving tear-film thinning from sparse time-series observations. Starting from reduced-order ODE models, we formulate a parameter inference framework that incorporates priors and identifiability considerations. To improve robustness and flexibility, we integrate neural network surrogates that learn corrections to the reduced models while respecting physical constraints. Synthetic and experimental data demonstrate that the hybrid ODE and NN approach outperforms purely mechanistic or purely data-driven methods. The results yield interpretable parameter estimates and improved predictions of break-up time, with potential clinical applications in understanding and diagnosing dry eye disease.

Alex Sheranko, University of Maryland, Baltimore County

A kernel-based SIR model for infectious spread in space and time

In this study, we consider an integro differential model describing the spread of infectious diseases in space and time, combining an SIR framework with spatial kernel interactions that capture nonlocal transmission. We derive the model, discuss well posedness conditions, and develop numerical methods for efficient simulation in large spatial domains. The kernel formulation accommodates heterogeneous contact patterns and mobility. We explore parameter identification from incidence data and examine how kernel choice affects wave speed and spatial clustering of infections. Applications include retrospective analysis of outbreaks and prospective scenario testing for intervention strategies such as localized quarantines and mobility restrictions. Numerical experiments illustrate the qualitative behavior of the model under different scenarios.

Machine Learning (4:20pm–5:00pm)

Noah Amsel, New York University

The polar express: Optimal matrix sign methods and their application to the Muon algorithm

Computing the polar decomposition and the related matrix sign function is a core primitive in numerical linear algebra and machine learning. We present new optimal order matrix sign iterations that minimize asymptotic error constants while maintaining numerical stability. Our analysis combines rational approximation theory with backward error considerations to derive iterations that outperform standard Newton and Zolotarev based schemes in practice. We demonstrate a key application to the Muon Algorithm, a recently proposed method in scientific machine learning, where improved matrix sign computations lead to faster convergence and better conditioning. Extensive experiments on synthetic and application driven problems show significant reductions in iteration counts and runtime, with our methods outperforming recent alternatives across a range of learning rates.

Leonardo Ferreira Guilhoto, University of Pennsylvania

Kolmogorov superposition meets physics-informed neural networks: A new architecture for scientific machine learning

Physics Informed Neural Networks (PINNs) have emerged as a flexible tool for solving PDE constrained inverse problems by embedding differential equation constraints into the loss. We introduce new architecture based on Kolmogorov's superposition theorem that decomposes multivariate mappings into compositions of univariate functions and linear combinations. The resulting model is more expressive per parameter and admits principled regularization aligning with PDE structure. We present training strategies that enforce physics constraints while stabilizing optimization, including curriculum schedules and adaptive residual weighting. Benchmarks on canonical PDEs demonstrate improved accuracy and data efficiency compared to standard PINNs, with faster convergence and better generalization.

A preprint of the full paper can be seen at: https://openreview.net/pdf?id=SyVPiehSbg

List of Presenters

Sylvia Amihere ~ University of Maryland, Baltimore County ~ samihere@umbc.edu
Noah Amsel ~ NYU ~ noah.amsel@nyu.edu
Ahmet Kaan Aydin ~ University of Maryland, Baltimore County ~ aaydin1@umbc.edu
Isaac Anthony Castro ~ University of Delaware ~ castroi@udel.edu
Leonardo Ferreira Guilhoto ~ University of Pennsylvania ~ guilhoto@sas.upenn.edu
Yifan Hu ~ University of Maryland, Baltimore County ~ yifan.hu@umbc.edu
Mikhail Lepilov ~ Rensselaer Polytechnic Institute ~ lepilm@rpi.edu
Jackie Lok ~ Princeton University ~ jackie.lok@princeton.edu
Arnab Roy ~ University of Delaware ~ arnabroy@udel.edu
Alex Sheranko ~ University of Maryland, Baltimore County ~ asheran1@umbc.edu
Shambhavi Suryanarayanan ~ Princeton University ~ ss3472@princeton.edu

Organizers

Benjamin Seibold, Temple University, benjamin.seibold@temple.edu

Daniel Szyld, Temple University, szyld@temple.edu