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 15, 2024 | 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 15, 2024)

9:15am-9:50am	Registration and Breakfast (provided)
9:50am-10:00am	Opening Remarks
10:00am-11:00am	**Presentations** (Algorithms)
11:00am-11:15am	*Coffee Break*
11:15am-12:15pm	**Presentations** (Data science)
12:15pm-1:30pm	Posters and Lunch (provided)
1:30pm-2:30pm	**Keynote Lecture** (Mark Embree)
2:30pm-2:45pm	*Coffee Break*
2:45pm-3:45pm	**Presentations** (Numerical Differential Equations)
3:45pm-4:00pm	*Coffee Break*
4:00pm-5:00pm	**Presentations** (Numerical Linear Algebra)
5:00pm-5:10pm	Closing Remarks
6:00pm-8:00pm	Group Dinner (attendance optional)

Speakers

Algorithms (10:00am-11:00am)

CHENYANG CAO, PURDUE UNIVERSITY

Kernel matrix approximations by sums of exponentials and stability of fast structured transforms

RISHAD ISLAM, LEHIGH UNIVERSITY

Acceleration of approximate maps for matrices arising in discretized PDEs

HAI ZHU, FLATIRON INSTITUTE

Recursive reduction quadrature for the evaluation of Laplace layer potentials in three dimensions

Data Science (11:15am-12:15pm)

MIKHAIL LEPILOV, EMORY UNIVERSITY

Estimating kernel matrix eigenvalues

SONIA REILLY, NEW YORK UNIVERSITY

Bayesian inversion of PDE-based problems using integrated nested Laplace approximations

ANNAN YU, CORNELL UNIVERSITY

Training an LTI system without an objective: A numerical analyst's perspectives on state-space models

Posters (12:15pm-1:30pm)

PIERRE AMENOAGBADJI, APAM COLUMBIA

Wave propagation in junctions of periodic half-spaces

CAROLINE HUBER, NEW YORK UNIVERSITY

Preconditioning without a preconditioner: Faster ridge-regression and Gaussian sampling with randomized block Krylov methods

SEAN T. MCQUADE, TEMPLE UNIVERSITY

Modeling lipid metabolism for multiple classes of virtual patient

AFRINA ASAD MEGHLA, TEMPLE UNIVERSITY

Modeling and simulation of calcium influx through NMDA receptors and its activation of TRPM4 channels

MADISON SHORAKA, TEMPLE UNIVERSITY

Modeling growth of pCF10-induced complex structures in Enterococcus faecalis biofilm under erythromycin treatment

JOVAN ZIGIC, MCMASTER UNIVERSITY

Stochastic variants of gradient descent for PDE-constrained optimization on Riemannian manifolds

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

MARK EMBREE, VIRGINIA TECH

Spectral Computations for Quasicrystals

Numerical Differential Equations (2:45pm-3:45pm)

HAMAD EL KAHZA, UNIVERSITY OF DELAWARE

Adaptive-rank implicit time integrator for advection-diffusion transport equations with inhomogeneous coefficients

ZACHARY MIKSIS, TEMPLE UNIVERSITY

A new Fick-Jacobs derivation with applications to computational branched diffusion networks

MANSUR SHAKIPOV, UNIVERSITY OF MARYLAND, COLLEGE PARK

Inf-sup stability of parabolic TraceFEM

Numerical Linear Algebra (4:00pm-5:00pm)

NOAH AMSEL, NEW YORK UNIVERSITY

Fixed-sparsity matrix approximation from matrix-vector products

ROBIN JOHN ARMSTRONG, CORNELL UNIVERSITY

"Collect, commit, expand": A strategy for faster CPQR-based column selection on short, wide matrices

TYLER CHEN, NEW YORK UNIVERSITY

Near-optimal hierarchical matrix approximation from matrix-vector products

Abstracts of Talks

Algorithms (10:00am-11:00am)

Chenyang Cao, Purdue University

Kernel matrix approximations by sums of exponentials and stability of fast structured transforms

In this talk, we show that some matrix transforms can be efficiently performed using sums-of-exponentials (SoE) approximations. Some existing transform algorithms based on SoE expansions are essentially matrix-vector multiplications with sequentially semiseparable (SSS) structures. The stability of the algorithms relating to such structures has been questioned. The study reveals two important factors that impact the stability of a structured algorithm: algorithm architecture and norm bounds of generators. The general SSS transforms are potentially unstable even if the relevant generators have norm bounds, while hierarchical structured algorithms are stable if the generators have bounded norms. SoE approximation is an effective way to control the norm of generators.

Rishad Islam, Lehigh University

Acceleration of approximate maps for matrices arising in discretized PDEs

Generally, discretization of partial differential equations (PDEs) creates a sequence of linear systems $A_k x_k = b_k$, k=1,2,...,N with well-known and structured sparsity patterns. For solving closely related systems in sequence, we can use preconditioner updates such as Sparse Approximate Map (SAM) instead of computing a preconditioner for each system from scratch. A SAM acts as a map from one matrix in the sequence to another nearby one for which we have an effective preconditioner. We seek to compute an optimal sparsity pattern to efficiently compute an effective SAM update. In this talk, we examine several sparsity patterns for computing the SAM update in an effort to characterize optimal or near-optimal sparsity patterns for linear systems arising from discretized PDEs. The allowable number of nonzeros in the sparsity pattern should strike a balance between the accuracy of the map and the cost to apply it in the iterative solver. We can show that the sparsity pattern of the exact map is a subset of the sparsity pattern of the transitive closure of a graph representation of A_k, $G(A_k)$. Additionally, we make use of the heterogeneous computing environment to accelerate the computation of the SAM. The inherently parallel nature of the SAM algorithm naturally lends itself towards efficient implementation in GPU and distributed computing systems. We present preliminary results in this area.

Hai Zhu, Flatiron Institute

Recursive reduction quadrature for the evaluation of Laplace layer potentials in three dimensions

We present an improved high-order nearly-singular layer potential evaluation scheme on curved patches. This method utilizes a combination of quaternionic harmonic polynomial approximation/extension of the layer potential density, along with a complete reduction from nearly-singular surface integrals to a set of endpoint evaluations at the vertices of the patch. Unlike its predecessor, this improved quadrature scheme relies on a cleaner representation of the intermediate line integrals derived from applying Stokes theorem on manifolds, and provides a more robust solution especially for targets located in close proximity to the patch boundary.

Data Science (11:15am-12:15pm)

Mikhail Lepilov, Emory University

Estimating kernel matrix eigenvalues

In data science, individual observations are often assumed to come independently from an underlying probability space. Kernel matrices formed from large sets of such observations arise frequently, for example during classification tasks. It is desirable to know the eigenvalue decay properties of these matrices without explicitly forming them, such as when determining if a low-rank approximation is feasible. In this talk, I introduce a new eigenvalue quantile estimation framework for some kernel matrices. This framework gives meaningful bounds for all the eigenvalues of a kernel matrix while avoiding the cost of constructing the full matrix. The kernel matrices under consideration come from a kernel with quick decay away from the diagonal applied to uniformly distributed sets of points in Euclidean space of any dimension. We prove the efficacy of this framework given certain bounds on the kernel function, and we provide empirical evidence for its accuracy. In the process, we also prove a very general interlacing-type theorem for finite sets of numbers. Additionally, we indicate an application of this framework to the study of the intrinsic dimension of data.

SONIA REILLY, NEW YORK UNIVERSITY

Bayesian inversion of PDE-based problems using integrated nested Laplace approximations

The core computational bottleneck in Bayesian inference is the cost of evaluating high-dimensional integrals. Often these integrals are computed using sampling-based methods such as MCMC, or by approximating the posterior with more tractable distributions through variational inference. A more recent alternative for problems which can be expressed as latent Gaussian models is known as Integrated Nested Laplace Approximations (INLA). INLA has been successfully applied to a variety of spatial statistics problems, but has not yet been applied to inverse problems derived from PDE-governed systems, which pose additional difficulties. In this presentation we will demonstrate a proof-of-concept application of INLA to one such PDE-based inverse problem.

ANNAN YU, CORNELL UNIVERSITY

Training an LTI system without an objective: A numerical analyst's perspectives on state-space models

State-space models (SSMs) leverage linear, time-invariant (LTI) systems to achieve state-of-the-art performance on tasks that involve long-range time-series and sequential data. A deep SSM utilizes thousands of LTI systems to solve a single task. Hence, there lacks a notion of "ground truth" for every single LTI system in an SSM. In this talk, we aim to answer the following question by studying the transfer functions of LTI systems: how do we know if an LTI system in an SSM is good or not without having an objective? By analyzing the initialization and training of the LTI systems, we propose mechanisms to improve the expressiveness, efficiency, and robustness of SSMs. Our method involves useful tools from numerical analysis, including the pseudospectral theory, singular value decomposition, and Fourier analysis.

Posters (12:15pm-1:30pm)

PIERRE AMENOAGBADJI, APAM COLUMBIA

Wave propagation in junctions of periodic half-spaces

This poster is devoted the 2D Helmholtz equation in presence of a periodic half-space. A numerical method has been proposed by Fliss, Cassan, Bernier (2010) to solve this equation under the critical assumption that the medium stays periodic in the direction of the interface. In fact, in this case, a Floquet-Bloch transform can be applied with respect to the variable along the interface, thus leading to a family of closed waveguide problems.

The purpose of this work is to deal with the case where the medium is no longer periodic in the direction of the interface, that is, if the periodic half-space is not cut in a direction of periodicity. As it is done by Gérard-Varet, Masmoudi (2015), we use the crucial (but non-obvious) observation that the medium has a quasiperiodic structure along the interface, namely, it is the restriction of a higher dimensional periodic structure. Accordingly, the idea is to interpret the studied PDE as the "restriction" of an augmented PDE in higher dimensions, where periodicity along the interface is recovered. This so-called lifting approach allows one to extend the ideas by Fliss, Cassan, Bernier (2010), but comes with the price that the augmented equation is non-elliptic (in the sense of the principal part of the differential operator), and thus more complicated to analyse and to solve numerically.

CAROLINE HUBER, NEW YORK UNIVERSITY

Preconditioning without a preconditioner: Faster ridge-regression and Gaussian sampling with randomized block Krylov methods

The preconditioned conjugate gradient algorithm is one of the most widely used methods for solving a positive definite linear system of equations $Ax = b$. We prove that by augmenting b with a random matrix, the block conjugate gradient algorithm outperforms preconditioned conjugate gradient (with a certain class of commonly used preconditioners). This provides some of the first theoretical evidence of the virtues of block Krylov subspace methods for solving a single linear system of equations. We also describe how our approach yields state of the art algorithms for computing the entire ridge regression regularization path and sampling Gaussian vectors.

SEAN T. MCQUADE, TEMPLE UNIVERSITY

Modeling lipid metabolism for multiple classes of virtual patient

Finding new pathways to effectively test drugs is extensive process which begins with modeling the biochemical pathway of interest. I present a framework to take a metabolic pathway and systematically generate virtual patients to simulate subpopulations. The subpopulations have parameters that allow their responses to simulated drug treatment to mimic those from earlier clinical trials.

AFRINA ASAD MEGHLA, TEMPLE UNIVERSITY

Modeling and simulation of calcium influx through NMDA receptors and its activation of TRPM4 channels

Calcium signaling plays an important role in intracellular communication, where calcium ions (Ca2+) act as a key step in signal transduction by regulating various cellular processes. In neurons, Ca2+ is important for synchronization of neuronal electrical activity, thereby supporting essential neuronal functions. Here we address the problem of Ca2+ entry into cells through NMDA receptors which subsequently activate the TRPM4 channels. This activation results in calcium binding to TRPM4, potentially initiating a neuron death signal and contributing to neurotoxicity. Our research focuses on whether the separation between NMDA receptors and TRPM4 channels could reduce NMDA-calcium-induced death signaling. By modeling calcium influx through NMDA receptors at varying distances from TRPM4, we aim to understand how this spatial separation might mitigate the toxic effects and could decrease excitotoxic cell death, which is particularly relevant in neurodegenerative diseases. To accurately represent calcium flux through individual GluN1/2A or GluN1/2B channels, we model both the time- and voltage-dependent channel conductance (g) and open probability (PO). To simulate the calcium influx, we applied a 2D diffusion equation with an NMDA receptor as the Ca2+ source, using specified boundary conditions with zero Neumann boundary conditions elsewhere. MATLAB was employed to simulate these time-dependent equations, allowing us to track calcium concentrations over time and space. Preliminary results indicate an 8% reduction in local calcium concentration with a 2 μm separation, suggesting a potential protective mechanism against excitotoxicity.

MADISON SHORAKA, TEMPLE UNIVERSITY

Modeling growth of pCF10-induced complex structures in Enterococcus faecalis biofilm under erythromycin treatment

The goal of this study is to construct a model that helps pinpoint characteristics of E. faecalis biofilms containing the pheromone responsive plasmid pCF10 that allow for its continued survival under erythromycin treatment. We employ the framework of stochastic cellular automata, and we utilize image processing techniques to assess the validity of the model and its assumptions.

JOVAN ZIGIC, MCMASTER UNIVERSITY

Stochastic variants of gradient descent for PDE-constrained optimization on Riemannian manifolds

The goal is to find a stochastic alternative to the Riemannian conjugate gradient (RCG) method for solving a PDE-constrained optimization problem. By defining Riemannian Stochastic Weighted (RSW) gradient methods, numerical studies suggest that an RSW variant is a suitable alternative to the RCG method.

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

MARK EMBREE, VIRGINIA TECH

Spectral Computations for Quasicrystals

In 2011, Dan Shechtman was awarded the Nobel Prize in Chemistry for his discovery of quasicrystals, novel materials with properties somewhere between the regularity of crystals and the disorder of random structures. In parallel with this scientific breakthrough, mathematicians have developed models of aperiodic order, such as Fibonacci substitutions and Penrose tilings. The self-adjoint linear operators based on such models often exhibit intriguing spectral structure. For example, the spectrum of the Fibonacci Hamiltonian is a zero-measure Cantor set. How can one approximate such fine structure using eigenvalue computations with finite dimensional matrices? Can one estimate quantities like fractal dimension of the spectrum and the integrated density of states? We will describe several aperiodic models, discuss their spectral properties, illustrate some numerical tools we can use to approach these problems, and show results from our calculations. This talk describes collaborations with Matt Colbrook, David Damanik, Jake Fillman, Anton Gorodetski, and May Mei.

Numerical Differential Equations (2:45pm-3:45pm)

HAMAD EL KAHZA, UNIVERSITY OF DELAWARE

Adaptive-rank implicit time integrator for advection-diffusion transport equations with inhomogeneous coefficients

In this work, we propose a high-order adaptive-rank implicit time integrator for time-dependent advection-diffusion PDEs with variable advection and diffusion coefficients. The discretized system takes the form of a Generalized Sylvester Equation, where state-of-the-art adaptive-rank solvers currently achieve a computational complexity of $O(N^3)$, where N is the resolution in a single dimension. In this talk, I will discuss our developed algorithm that leverages two key strategies: (i) constructing solution bases from Krylov-like subspaces, and (ii) preconditioning the solve for the coefficient matrix that composes these basis vectors. Our preconditioning strategy scales as $O(r^3)$, where r is the rank, reducing the overall computation to a super-optimal complexity of O(N) when the solution exhibits a low-rank structure. We validate the algorithm's potential through a series of numerical examples, including cases with equilibrated advection-diffusion and advection-dominated ones. This algorithm forms a critical building block for scalable and efficient solvers for high-dimensional, nonlinear time-dependent problems.

ZACHARY MIKSIS, TEMPLE UNIVERSITY

A new Fick-Jacobs derivation with applications to computational branched diffusion networks

The Fick-Jacobs equation is a classical model reduction of 3-dimensional diffusion in a tube of varying radius to a 1-dimensional problem with radially scaled derivatives. This model has been shown to be unstable when the radial gradient is too steep. In this work, we present a new derivation of the Fick-Jacobs equation that results in the addition of higher order spatial derivative terms that provide additional stability in a wide variety of cases and improved solution convergence. We also derive new numerical schemes for branched nodes within networks and provide stability conditions for these schemes. The computational accuracy, efficiency, and stability of our method is demonstrated through a variety of numerical examples.

MANSUR SHAKIPOV, UNIVERSITY OF MARYLAND, COLLEGE PARK

Inf-sup stability of parabolic TraceFEM

We develop a parabolic inf-sup theory for a modified TraceFEM semi-discretization in space of the heat equation posed on a stationary surface embedded in R^2. We consider the normal derivative volume stabilization and add an L^2-type stabilization to the time derivative. We assume that the representation of and the integration over the surface are exact, however, all our results are independent of how the surface cuts the bulk mesh. For any mesh for which the method is well-defined, we establish necessary and sufficient conditions for inf-sup stability of the proposed TraceFEM in terms of H^1-stability of a stabilized L^2-projection and of an inverse inequality constant that accounts for the lack of conformity of TraceFEM. Furthermore, we prove that the latter two quantities are bounded uniformly for a sequence of shape-regular and quasi-uniform bulk meshes. We derive several consequences of uniform discrete inf-sup stability, namely uniform well-posedness, discrete maximal parabolic regularity, parabolic quasi-best approximation, convergence to minimal regularity solutions, and optimal order-regularity energy and L^2-error estimates. We show that the additional stabilization of the time derivative restores optimal conditioning of time-discrete TraceFEM typical of fitted discretizations.

Numerical Linear Algebra (4:00pm-5:00pm)

NOAH AMSEL, NEW YORK UNIVERSITY

Fixed-sparsity matrix approximation from matrix-vector products

We study the problem of approximating a matrix A with a matrix that has a fixed sparsity pattern (e.g., diagonal, banded, etc.), when A is accessed only by matrix-vector products. We describe a simple randomized algorithm that returns an approximation with the given sparsity pattern with Frobenius-norm error at most $(1+\varepsilon)$ times the best possible error. When each row of the desired sparsity pattern has at most s nonzero entries, this algorithm requires $O(s/\varepsilon)$ non-adaptive matrix-vector products with A. We also prove a matching lower-bound, showing that, for any sparsity pattern with Theta(s) nonzeros per row and column, any algorithm achieving $(1+\varepsilon)$ approximation requires $\Omega(s/\varepsilon)$ matrix-vector products in the worst case. We thus resolve the matrix-vector product query complexity of the problem up to constant factors, even for the well-studied case of diagonal approximation, for which no previous lower bounds were known.

ROBIN JOHN ARMSTRONG, CORNELL UNIVERSITY

"Collect, commit, expand": A strategy for faster CPQR-based column selection on short, wide matrices

Column-pivoted QR (CPQR) factorization is a computational primitive used in numerous applications that require selecting a small set of "representative" columns from a much larger matrix. These include applications in kernel-based machine learning and low-rank approximation, where the matrices in question often have a moderate number of rows but an extremely large number of columns. This talk introduces a modification of the Golub-Businger algorithm which, for many matrices of this type, can perform CPQR-based column selection much more efficiently. This algorithm, which we call CCEQR, is based on a three-step "collect, commit, expand" strategy that limits the number of columns being manipulated, while also transferring more computational effort from level-2 BLAS to level-3. Unlike most CPQR algorithms that exploit level-3 BLAS, CCEQR is deterministic, and provably recovers a column permutation equivalent to the one computed by the Golub-Businger algorithm. Through tests on spectral clustering and low-rank approximation problems, we will demonstrate that CCEQR can significantly outperform GEQP3.

TYLER CHEN, NEW YORK UNIVERSITY

Near-optimal hierarchical matrix approximation from matrix-vector products

We describe a randomized algorithm for producing a near-optimal hierarchical off-diagonal low-rank (HODLR) approximation to an nxn matrix $\mathbf{A}$, accessible only though matrix-vector products with $\mathbf{A}$ and $\mathbf{A}^T$ Our algorithm can be viewed as a robust version of widely used "peeling" methods for recovering HODLR matrices and is, to the best of our knowledge, the first matrix-vector query algorithm to enjoy theoretical worst-case guarantees for approximation by any hierarchical matrix class. To control the propagation of error between levels of hierarchical approximation, we introduce a new perturbation bound for low-rank approximation, which shows that the widely used Generalized Nyström method enjoys inherent stability when implemented with noisy matrix-vector products. We also introduced a novel randomly perforated matrix sketching method to further control the error in the peeling algorithm.

List of Participants

Pierre Amenoagbadji, *APAM Columbia*, ka3012@columbia.edu

Noah Amsel, *New York University*, noah.amsel@nyu.edu

Robin J. Armstrong, *Center for Applied Mathematics, Cornell University*, rja243@cornell.edu

Gonzalo A. Benavides, *University of Maryland, College Park*, gonzalob@umd.edu

Mingchao Cai, *Morgan State University*, Mingchao.Cai@morgan.edu

Chenyang Cao, *Purdue University*, cao302@purdue.edu

Arielle Carr, *Lehigh University*, arg318@lehigh.edu

Tyler Chen, *New York University*, tyler.chen@nyu.edu

Gitansh Dandona, *University of Delaware*, gitdan@udel.edu

Hamad El Kahza, *University of Delaware*, helkahza@udel.edu

Caroline Huber, *New York University,* ch4152@nyu.edu

Rishad Islam, *Lehigh University*, mds222@lehigh.edu

Dylan T. Jacobs, *Swarthmore College*, djacobs2@swarthmore.edu

Mikhail Lepilov, *Emory University*, mlepilo@emory.edu

Zachary Miksis, *Temple University*, miksis@temple.edu

Joseph Nakao, *Swarthmore College*, jnakao1@swarthmore.edu

Blessing U. Nwonu, *Temple University*, blessing.nwonu@temple.edu

Jingmei Qiu, *University of Delaware*, jingqiu@udel.edu

Sonia M. Reilly, *Courant Institute, NYU,* sonia.reilly@nyu.edu

Francois-Henry Rouet, *Ansys*, francois-henry.rouet@ansys.com

Mansur Shakipov, *University of Maryland, College Park*, shakipov@umd.edu

Rick Smith, *Unaffiliated*, rsmithOIT@yahoo.com

Annan Yu, *Cornell University*, ay262@cornell.edu

Hai Zhu, *Flatiron Institute,* hzhu@flatironinstitute.org

Organizers

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

Daniel Szyld, Temple University, szyld@temple.edu