27th Conference of the International Linear Algebra Society (ILAS 2026)

Virginia Tech, Blacksburg, VA, USA

May 18–22, 2026

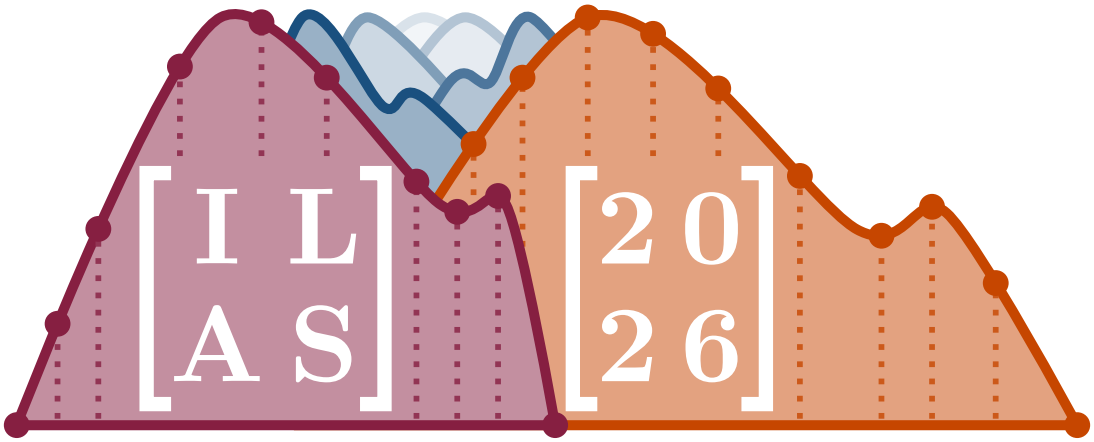

Book of Abstracts

Table of Contents

Introduction

The 27th Conference of the International Linear Algebra Society (ILAS 2026) is held on the campus of the Virginia Polytechnic Institute and State University (Virginia Tech) in Blacksburg, Virginia, from May 18 to 22, 2026. The conference theme, *"Linear Algebra on the Blue Ridge: Panoramas of Theory and Application,"* reflects both the scenic mountain setting of southwest Virginia and the broad scope of the meeting, which brings together researchers from across the full spectrum of linear algebra from core theory and numerical analysis to applications and education.

The conference is organized by a Local Organizing Committee from Virginia Tech, with the scientific program decided by an international Scientific Committee drawing on expertise from institutions across Europe, North America, and Asia.

ILAS 2026 is an official satellite conference of the International Congress of Mathematicians 2026, to be held in Philadelphia, PA, July 23–30, 2026.

Content

About ILAS

Founded in 1989, the International Linear Algebra Society (ILAS) is a worldwide community of mathematicians, researchers, and educators united by their commitment to advancing linear algebra and its applications. The Society supports its members through a range of activities: it publishes the Electronic Journal of Linear Algebra (ELA), a platinum open-access journal dedicated to high-quality research in matrix analysis and linear algebra, as well as IMAGE, a semiannual bulletin featuring news, interviews, and reports from across the community. Members stay connected through ILAS-Net, a moderated electronic newsletter covering conferences, announcements,

and developments in the field. ILAS awards several internationally recognized prizes to linear algebra researchers at different stages in their careers, including the Hans Schneider Prize, the ILAS Taussky-Todd Prize, and the ILAS Richard A. Brualdi Early Career Prize.

Conference History

ILAS conferences have taken place since 1989 over four continents, North America, Europe, Asia, and South America, featuring high-profile speakers presenting the latest developments in linear algebra research and its wide range of applications. The complete list of ILAS conferences is as follows:

#	Location	Dates
1	Provo, Utah, USA	August 12–15, 1989
2	Lisbon, Portugal	August 3–7, 1992
3	Pensacola, Florida, USA	March 17–20, 1993
4	Rotterdam, The Netherlands	August 15–19, 1994
5	Atlanta, Georgia, USA	August 16–19, 1995
6	Chemnitz, Germany	August 14–17, 1996
7	Madison, Wisconsin, USA	June 3–6, 1998
8	Barcelona, Spain	July 19–22, 1999
9	Haifa, Israel	June 25–29, 2001
10	Auburn, Alabama, USA	June 10–13, 2002
11	Coimbra, Portugal	July 19–22, 2004
12	Regina, Saskatchewan, Canada	June 26–29, 2005
13	Amsterdam, The Netherlands	July 18–21, 2006
14	Shanghai, China	July 16–20, 2007
15	Cancun, Mexico	June 16–20, 2008
16	Pisa, Italy	June 21–25, 2010
17	Braunschweig, Germany	August 22–26, 2011
18	Providence, Rhode Island, USA	June 3–7, 2013
19	Seoul, Korea	August 6–9, 2014
20	Leuven, Belgium	July 11–15, 2016
21	Ames, Iowa, USA	July 24–28, 2017
22	Rio de Janeiro, Brazil	July 8–12, 2019
23	Virtual meeting	May 17–21, 2021
24	Galway, Ireland	June 20–24, 2022
25	Madrid, Spain	June 12–16, 2023
26	Kaohsiung, Taiwan	June 23–27, 2025
27	Blacksburg, Virginia, USA	May 18–22, 2026

Local Organizing Committee

Christopher Beattie (Virginia Tech)
Paul Cazeaux (Virginia Tech)
Eric de Sturler (Virginia Tech)
Mark Embree (Virginia Tech)
Serkan Gugercin (Virginia Tech)
Agnieszka Miedlar (Virginia Tech)
Mirjeta Pasha (Virginia Tech)
Megan Wawro (Virginia Tech)
Steffen Werner (Virginia Tech)

Scientific Committee

Shreemayee Bora (IIT Guwahati)
Geir Dahl (University of Oslo)
Eric de Sturler (Virginia Tech)
Ioana Dumitriu (University of California San Diego)
Stefan Güttel (University of Manchester)
Misha Kilmer (Tufts University)
Agnieszka Miedlar (Virginia Tech)
André Ran (Vrije Universiteit Amsterdam)
Helena Smigoc (University College Dublin)
David Strong (Pepperdine University)
Daniel Szyld (Temple University)
Raf Vandebril (KU Leuven)
Heather Wilber (University of Washington)

Special Sessions

Mentoring Panel (Sponsored by the National Science Foundation (NSF))

Date, Time & Venue: Tuesday, May 19, 2026, 5:30pm – 6:30pm, McBryde Hall 100

A mentoring panel will bring together senior academics, national laboratory researchers, and industry experts to share their perspectives on careers and research across the various domains of linear algebra.

This panel aims to provide guidance and support to early-career researchers through a discussion of mentoring experiences, challenges, and best practices across different academic systems. This activity is part of the broader service component of an National Science Foundation Early Career Award, through which we have received support to fund participation of U.S.-based early-career conference speakers.

Panelists:

Daniela Calvetti (Case Western Reserve University)
Shaun Fallat (University of Regina)
Heike Faßbender (TU Braunschweig)
Misha Kilmer (Tufts University)
Francoise Tisseur (The University of Manchester)
Roel Van Beeumen (Lawrence Berkeley National Laboratory)

Moderator: Mirjeta Pasha (Virginia Tech)

ILAS Education Workshop: Bridging Abstract and Numerical Linear Algebra

Date, Time & Venue: Thursday, May 21, 2026, 2:00pm – 3:40pm, Torgersen Hall 2150

Workshop Organizers: Sepideh Stewart, Mike Michailidis, Rachel Quinlan

This workshop will discuss ways of bridging abstract and numerical linear algebra. Our topics are included, but not limited to, the following:

- How to introduce numerical linear algebra to students (some examples of coding and programming will be provided),
- The role of AI in teaching and learning linear algebra,
- The role of technology in teaching linear algebra,
- Collaboration, creating and sharing resources (GitHub), lesson plans, and publications

Everybody is welcome!

See the associated website at

https://ilas-ed.math.ou.edu/conferences-and-workshops/workshop2026/

ILAS Business Meeting

Date, Time & Venue: Thursday, May 21, 2026, 4:15 pm – 6:00 pm, McBryde Hall 100

Agenda:

A. Reading of Notice of Meeting.
B. Review of the Minutes of Previous Meeting.
C. Report of President and Vice President.
D. Report of Secretary/Treasurer.
E. Reports of Committees.
   i) Education Committee (chair: Sepideh Stewart).
   ii) Outreach and Membership Committee (chair: Shahla Nasserasr).

iii) Journals Committee (chair: Zlatko Drmac).

iv) Electronic Journal of Linear Algebra (Editor in Chief: Vanni Noferini).

v) IMAGE (Editor in Chief: Louis Deaett).

F. Any Other Business.

G. Adjournment.

LAA Special Issue for ILAS 2026

Linear Algebra and its Applications (LAA) is pleased to announce a special issue on the occasion of the 27th Conference of the International Linear Algebra Society (ILAS 2026). This special issue is part of the long-standing collaboration between Elsevier and the International Linear Algebra Society and continues the tradition of many previous special issues related to ILAS Conferences.

Papers based on talks presented at the 27th Conference of the International Linear Algebra Society, whose authors include at least one conference participant as a (co)author and which align with the aims and scope of LAA, should be submitted through the LAA submission system.

The submission opening date is May 25, 2026.
The submission deadline is October 31, 2026.

The special issue editors for this issue are:

Haim Avron (Tel Aviv University, Israel)
Geir Dahl (University of Oslo, Norway)
Agnieszka Miedlar (Virginia Tech, USA)
Hugo Woerdeman (Drexel University, USA)

In submitting a paper, please select `VSI:LAA_ILAS2026` as the article type and Froilán M. Dopico as the responsible editor-in-chief.

Authors may suggest one of the special issue editors above to handle their submission. Papers will be referred according to the usual standards of LAA.

Sponsors and Supporters

ILAS 2026 is generously supported by the following organizations:

University Sponsors:

- Virginia Tech
- College of Science, Virginia Tech
- Office of Research and Innovation, Virginia Tech
- Academy of Data Science, Virginia Tech
- Department of Mathematics, Virginia Tech

Professional Organizations:

- International Linear Algebra Society (ILAS)
- Society for Industrial and Applied Mathematics (SIAM)

Independent Federal Funding Agency:

- National Science Foundation (NSF)

Coorporate Sponsors:

- MathWorks
- Procedura

Publishers:

- Elsevier
- Taylor and Francis

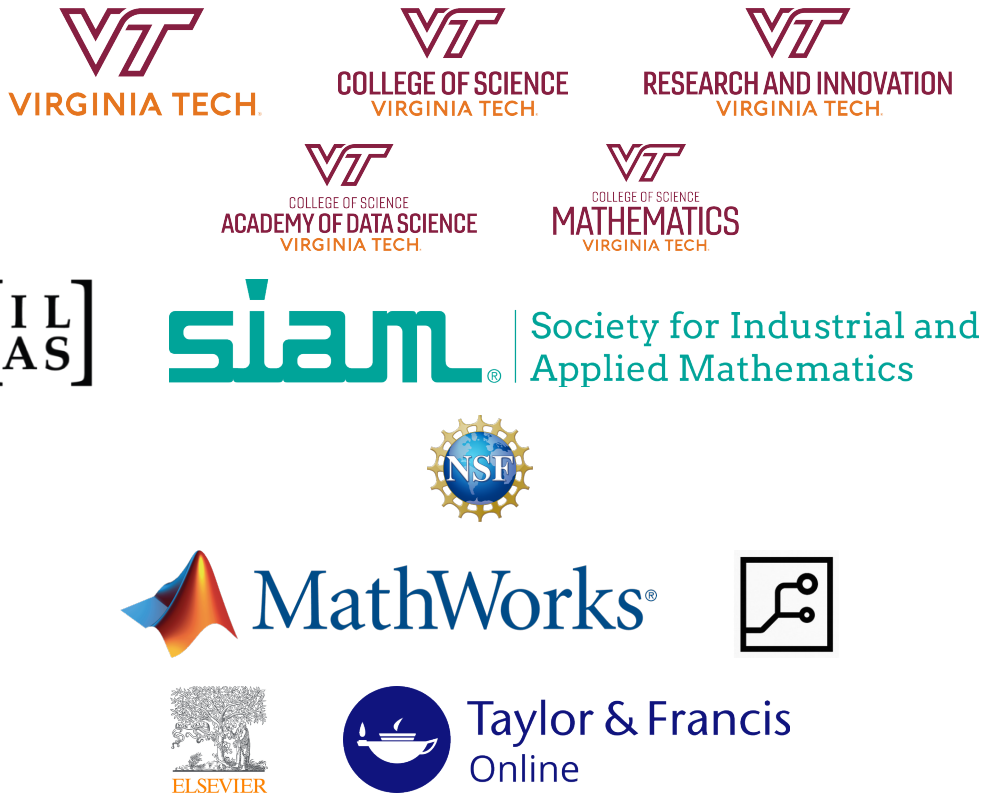

Overview Schedule

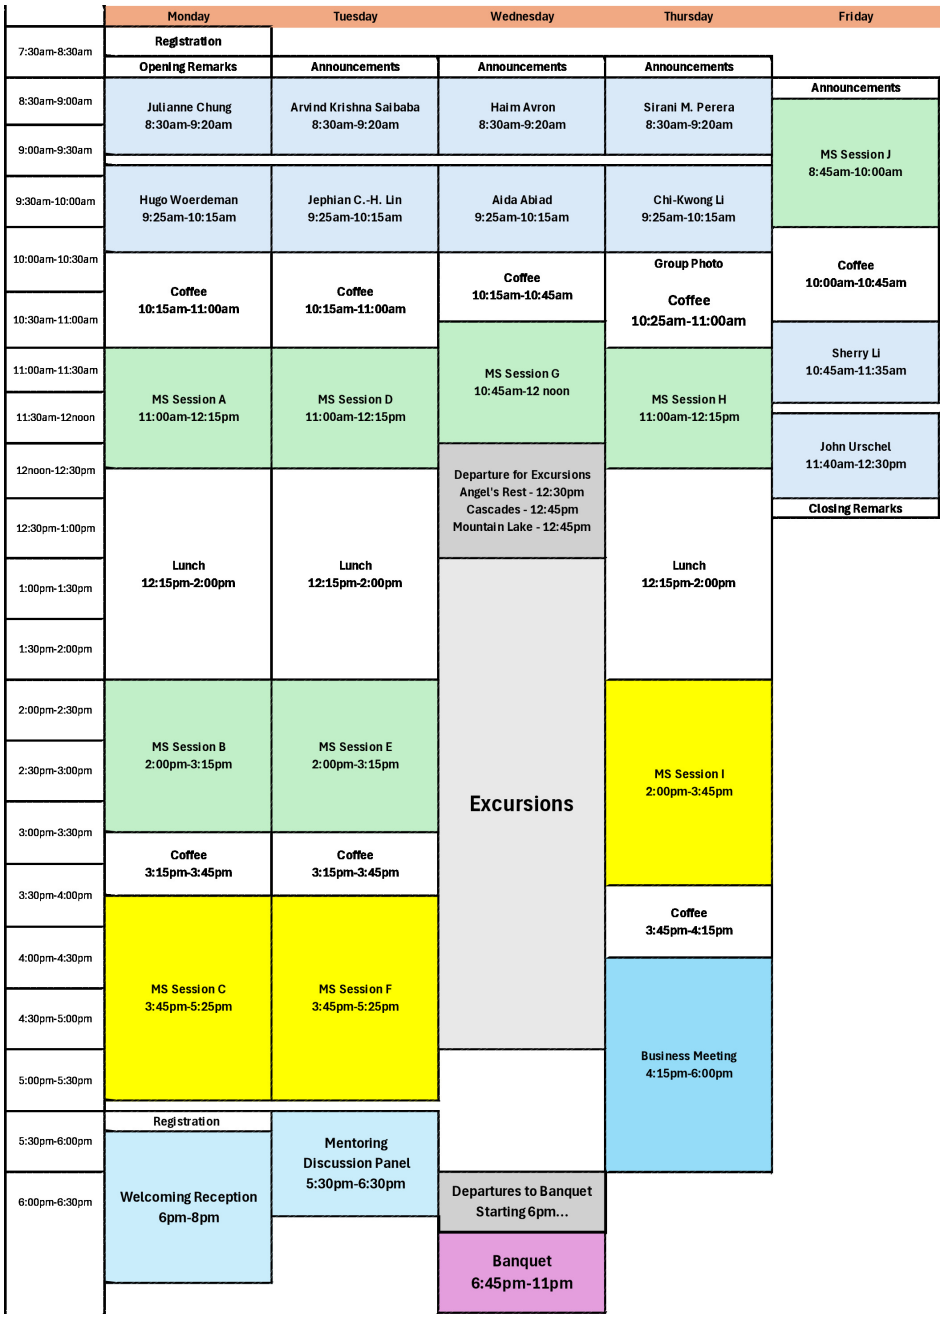

More details about the daily schedule can be seen in the program-at-a-glance

https://ilas2026.math.vt.edu/docs/program-at-a-glance.pdf

Local Information

Founded in 1872 as the land-grant Virginia Agricultural and Mechanical College, Virginia Tech is now a comprehensive, innovative research university with the largest full-time student population in Virginia. Through a combination of its three missions of learning, discovery, and engagement, Virginia Tech continually strives to accomplish the charge of its motto: Ut Prosim (That I May Serve).

At a glance . . .

- Located in Blacksburg, Virginia,
- Nine colleges and a graduate school,
- 110+ undergraduate programs,
- 120+ master's and doctoral degree programs,
- 38,000+ full-time students,
- 2,500+ instructional faculty,
- Includes an 2,600 acres main campus, an innovation campus in northern Virginia, the Health Sciences and Technology Campus in Roanoke, a study-abroad site in Switzerland and a 1,800 acres research farm,
- Adjacent corporate research center houses more than 220 companies.

The following overview map shows the location of the main ILAS 2026 conference buildings: McBryde Hall, Torgersen Hall, and Goodwin Hall.

Blacksburg Map
Conference buildings and near business areas highlighted

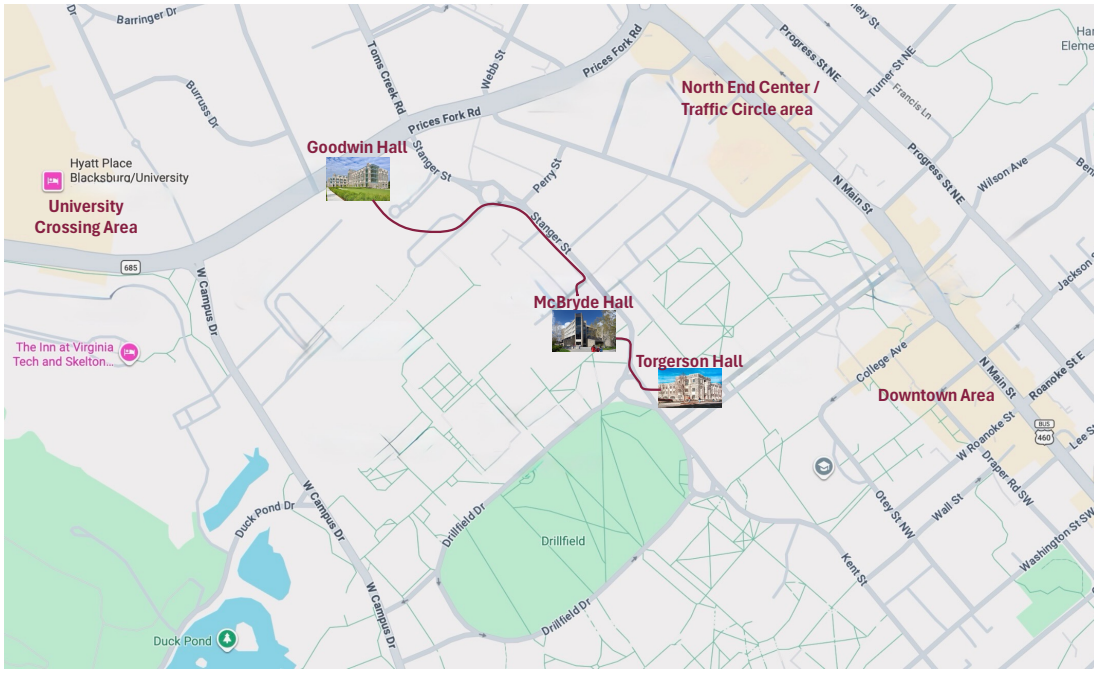

Plenary Talks

List of Talks

The Eigenvalue Method in coding theory

Author: Aida Abiad[1]

[1] *Eindhoven University of Techonolgy*

Contact: a.abiad.monge@tue.nl

One of the main goals in spectral graph theory is to deduce the principal properties and structure of a graph from its graph spectrum. In this talk we will show how spectral graph theory provides powerful methods for obtaining results concerning substructures of graphs, and also how these results can be useful in other mathematical fields such as coding theory. In particular, we will derive sharp eigenvalue bounds for the k-independence number of a graph (or equivalently, the independence number of the k-th graph power), which is known to be very hard to compute. We will see how to use polynomials and mixed integer linear programming in order to optimize and compute such spectral bounds. Finally, we will illustrate some recent applications of the obtained eigenvalue bounds to coding theory. The obtained results are encouraging and strongly suggest that spectral graph theory can uncover structural properties of ambient spaces that are relevant

to coding theory, but that are often not captured by classical coding theory techniques. There is no question that eigenvalues play a central role in our fundamental understanding of graphs. The goal of this talk is to show that we can also use them for deepening our understanding of codes.

Quantum Numerical Linear Algebra Doesn't Have to be Hard: A Matrix Algebra Oriented Approach

Author: Haim Avron[1]

[1] *Tel Aviv University*

Contact: haimav@tauex.tau.ac.il

The field of quantum computing offers a unique opportunity to revolutionize numerical linear algebra and scientific computing. This stems from the ability of quantum computers to efficiently model complex structures, and to represent and manipulate high-dimensional vectors and matrices using exponentially fewer qubits. These advantages arise from the fundamental principles of superposition and entanglement inherent to qubits.

Yet, the current landscape of quantum computing research is dominated by intricate, tailor-made circuit designs developed in an ad-hoc manner for specific mathematical challenges. Although such state-of-the-art quantum algorithms provide a powerful means of translating diverse computations into circuits, their development is far from straightforward. It often demands extensive "circuit engineering" to achieve desired mathematical outcomes. This also extends to quantum numerical linear algebra.

In this talk, I will discuss our recent progress on developing a unified and systemic approach to utilizing quantum computing for numerical linear algebra. Our research centers around a novel Quantum Linear Algebra (qLA) framework offering fundamental matrix algebra building blocks, akin to BLAS – but for Quantum Computers.

A new age of iterative Krylov methods for inverse problems: What to do with expensive inner-product computations?

Author: Julianne Chung[1]

[1] *Emory University*

Contact: jmchung@emory.edu

Iterative Krylov projection methods have become widely used for solving large-scale linear inverse problems. However, methods based on orthogonality include computations of inner-products, which becomes costly when the number of iterations is high, are a bottleneck for parallelization, and can cause the algorithms to break down due to information loss in the projections.

In this talk, I will describe two new approaches to handle expensive inner-product computations in the context of solving large-scale inverse problems. First, we describe new iterative solvers based on the randomized Golub-Kahan approach, where sketched inner products are used to estimate inner products of high-dimensional vectors. Second, we describe new inner-product-free Krylov solvers that avoid inner-products completely. For both approaches, we also describe hybrid methods that combine iterative projection methods with Tikhonov regularization, where regularization parameters can be selected automatically during the iterative process. Numerical results from image reconstruction show the potential benefits of these approaches.

Numerical Ranges and Dilations: Theoretical Advances and Applied Perspectives

Author: Chi-Kwong Li[1]

[1] *College of William and Mary*

Contact: ckli@math.wm.edu

The numerical range is a fundamental tool for understanding the properties of matrices and operators. In this talk, we discuss recent advances in the study of the numerical range and its generalizations, specifically focusing on their utility in analyzing operator dilations. We demonstrate how these theoretical frameworks provide critical insights into applied topics, including quantum information theory and modern data science.

Geometric and Algebraic Methods for Constructing Hierarchically Low-Rank Matrices

Author: Sherry Li[1]

[1] *LBNL*

Contact: xsli@lbl.gov

Hierarchically low-rank (H-LR) matrices have been widely used to design fast solvers for integral equations, boundary element methods, discretized PDEs, and kernel matrices in statistical and machine learning. The computational bottleneck in these solvers is often the construction algorithm which converts a standard dense matrix into an H-LR format. We will present two types of algorithms for fast construction of H-LR matrices. One type of algorithm exploits geometric information that works well with high dimensional data from ML applications. Another type of algorithm is purely algebraic, which works well for building general-purpose linear solvers. In both algorithms, randomization plays a significant role.

Inverse problems on a graph: strong matrices and graph minors

Author: Jephian C.-H. Lin[1]

[1] *National Yang Ming Chiao Tung University*

Contact: jephianlin@gmail.com

Inverse problems on a graph investigate how spectral behaviors interact with the matrices associated with the given graph. Such problems not only uncover structural information about the graph from its spectral data, but also identify fundamental properties shared by all matrices defined on the graph. A classic example is the Colin de Verdière parameter, which characterizes planarity via the maximum nullity over matrices defined on the graph. Central to these studies are strong properties, which either preserve the matrix pattern while perturbing spectral data or preserve the spectral data while adjusting the pattern. A recurring theme is that if a spectral behavior is realizable by a strong matrix on a graph, then it often remains realizable for any graph containing it as a minor. In this talk, we will survey these connections and present new results on strong properties for discrete Schrödinger operators.

When Structured Matrices Pay Off: Linear Algebra at the Heart of Wireless Communication

Author: Sirani M. Perera[1]

[1] *Embry-Riddle Aeronautical University, USA*

Contact: pereras2@erau.edu

Millimeter waves within the sub-terahertz band offer a plethora of applications in next-generation wireless communication. However, they also introduce severe real-time and hardware limitations, making conventional wideband multi-beam beamforming exceedingly complex. For example, an N-element array using true-time delay beamformers needs $\mathcal{O}(N^2)$ time delays or phase shifts, while FFT-based beamformers experience the beam-squint problem. Thus, this lecture addresses how structured matrix analysis can lead to algorithmic and architectural solutions to address these obstacles.

The central object of the talk is the delay Vandermonde matrix (DVM), a new structured matrix whose elements are defined via true-time delays in the spatio-temporal domain. By exploring its structure, we derive fast DVM algorithms that reduce computational and hardware complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ and also to $\mathcal{O}(N)$ for N beams. These results position DVM-based beamformers as a superclass of FFT-based multibeam beamformers. Unlike FFT beams, DVM beamformers are inherently wideband and eliminate beam squint, while simultaneously reducing chip area, power consumption, and system cost.

Beyond this, DVM factorization enhances AI-chip design by using structured weight matrices in neural network layers, enabling real-time wireless communication. This model-driven approach

constrains learning through a structured neural network, enabling real-time inference on massive data streams at gigahertz rates. Taken together, the lecture emphasizes that structured matrices connect algorithmic complexity, hardware feasibility, and future wireless communication powered by AI.

Stochastic trace estimation for parameter-dependent matrices

Author: Arvind Krishna Saibaba[1]

[1] *North Carolina State University*

Contact: asaibab@ncsu.edu

Estimating the trace of a matrix, that is only accessible by matrix-vector products, is a fundamental task in scientific computing and data science. This has many applications including network analysis, estimation of matrix norms and spectral densities, estimation of log-determinants, etc. Monte Carlo methods is one of the prevalent approaches for estimating the trace of the matrix. We consider Monte Carlo methods for minimizing the trace of a symmetric matrix $A(\theta)$, parameterized (perhaps, nonlinearly) by a parameter θ from a compact space $\Theta \subset \mathbb{R}^d$. We derive bounds for the number of samples required to estimate the backward error of the Monte Carlo to a given accuracy with high probability. The bounds are based on epsilon nets and chaining and use ideas from statistical learning theory and high dimensional probability. We illustrate this problem using applications to optimal experimental design and hyperparameter estimation in Bayesian inverse problems. Joint work with Ilse Ipsen.

Nodal Statistics for Graphs and Matrices

Authors: John Urschel[1], Dan Mikulincer[2], Lior Alon[1]

[1] *MIT*

[2] *University of Washington*

Contact: urschel@mit.edu

Given a symmetric matrix with a given sign pattern, what can the sign patterns of its eigenvectors look like? This simple question is closely related to the study of discrete nodal statistics, and draws strong parallels with classical results in analysis for Laplacian eigenfunctions. In this talk, we will give an overview of the field, covering key results on nodal sets for graphs and their connection to known results and open problems in the continuous setting. In addition, we will discuss some recent progress towards a more complete understanding of the extremal properties of the nodal statistics of a matrix.

Optimal interpolation in Hardy, Bergman and ℓ^p_A spaces: a reproducing kernel Banach space approach

Author: Hugo Woerdeman[1]

[1] *Drexel University*

Contact: hjw27@drexel.edu

After a review of the reproducing kernel Banach space framework and semi-inner products, we apply the techniques to the settings of sequence spaces ℓ^p (including the finite dimensional case), the associated function space ℓ^p_A, Hardy spaces H^p and Bergman spaces A^p, $1 < p < \infty$, on the unit ball in $\mathbb{C}^n$, as well as the Hardy space on the polydisk and half-space. In particular, we show how the framework leads to a procedure to find a minimal norm element f satisfying interpolation conditions $f(z_j) = w_j$, $j = 1, \ldots, n$. We also provide numerical examples.

This talk is based on joint works with Gilbert Groenewald, Sanne ter Horst and Eder Kikianty.

MS 01: Numerical Linear Algebra in Machine Learning

Minisymposium Organizers: Xinye Chen, EL-Mehdi El Arar

Numerical linear algebra is pivotal in advancing machine learning by enabling efficient and accurate computations for large-scale models. This minisymposium explores recent advances in numerical methods tailored for machine learning tasks, including optimization techniques, tensor factorizations, mixed precision analysis, and high-performance computing. Emphasis will be placed on algorithms that enhance computational efficiency, such as those using low-precision arithmetic, while ensuring numerical stability through rigorous error analysis. This minisymposium will bring together experts to discuss novel approaches, theoretical insights, and practical implementations that address challenges in training/inference and deploying machine learning algorithms, with a focus on numerical linear algebra techniques and their theoretical underpinnings through rounding error analysis. Held as part of the 27th Conference of the International Linear Algebra Society (ILAS 2026) at Virginia Tech, Blacksburg, VA, USA, from May 18–22, 2026, this event will foster vibrant discussions and collaborations among leading researchers in the field.

List of Talks

Look-ahead mixed-precision inference of LLMs

Author: Stanislav Budzinskiy[1]

[1] *University of Vienna*

Contact: stanislav.budzinskiy@univie.ac.at

We address the floating-point computation of compositionally-rich functions, concentrating on LLM inference. Based on the rounding error analysis of a composition, we provide an adaptive strategy to select components of the inner function that need to be recomputed more accurately to improve the numerical stability. We explain how this strategy can be applied to different compositions within a transformer neural network and illustrate its overall effect on LLM inference.

Understanding and Leveraging Adaptive Algorithms' Sensitivity to Change-of-Basis

Authors: Adela DePavia[1], Rebecca Willett[1], Vasileios Charisopoulos[2]

[1] *The University of Chicago*

[2] *The University of Washington*

Contact: adepavia@uchicago.edu

Adaptive gradient optimization algorithms—including Adam, Adagrad, and their variants—have found widespread use in machine learning, signal processing, and many other settings. However many algorithms in this family are not rotationally equivariant: in this talk we examine how a simple change-of-basis in either parameter space or data space can drastically impact both the convergence rates and the generalization of these algorithms. We begin by studying reparameterizations in parameter space, and describe a data-driven method proposed in our recent work which produces a "favorable" basis for adaptive algorithms. Our method is an orthonormal transformation based on the expected gradient outer product (EGOP) matrix. We present theoretical results and empirical evidence that reparameterizations based on the EGOP eigenbasis can improve convergence of adaptive gradient methods, even when these leading eigenspaces are approximated using randomized numerical linear algebra methods. We show that for a broad class of functions, the sensitivity of adaptive algorithms to choice-of-basis is influenced by the decay of the EGOP matrix spectrum. We illustrate the potential impact of EGOP reparameterization by presenting empirical evidence and theoretical arguments that common machine learning tasks with "natural" data exhibit EGOP spectral decay.

Universal Kronecker Core Factorization of the NTK: Quantifying Implicit Bias of Gradient Descent

Authors: Eli Shlizerman[1], Eric Shea-Brown[1], James Hazelden[1], Laura Driscoll[2]

[1] *University of Washington*

[2] *The Allen Institute*

Contact: jhazelde@uw.edu

How are learned representations incrementally formed to solve tasks by Gradient Descent (GD)? In this talk, we will show that each step of GD is exactly given by the application of a massive tensor-valued linear operator, which we call the Configuration Space Neural Tangent Kernel (NTK). We prove that it can be decomposed into two operators: P and K, the former capturing state-to-state dynamical dependencies and the latter capturing immediate parameter-to-state dependencies. Based on this, we prove a universal result stating that any weight-based model in deep-learning or optimal control can be factored so that K is a simpler Kronecker product matrix (dubbed the Kronecker core factorization). Importantly, the key ingredients in this core are already computed during inference, hence they can be used to immediately determine which tasks are more easily learnable by GD for a particular model. We show that this structure implies a bottlenecking of GD dynamics, leading to low-rank dynamical corrections and implicit bias towards particular tasks. We specialize our analysis to recurrent models (e.g., neural ODEs or RNNs), showing how the factorization constrains the temporal modes of learning. In addition to these theoretical results, we develop a package (kpflow) using matrix-free numerical linear algebra fast analysis of the NTK gradient operator for any model. This package makes it easier to work with linear operators on tensor domains with fast, randomized analysis tools.

Reduced Rank Extrapolation for Matrix Equations

Authors: Patrick Kürschner[1], Peter Benner[2], Xiaobo Liu[2]

[1] *Leipzig University of Applied Sciences*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems*

Contact: xliu@mpi-magdeburg.mpg.de

Reduced rank extrapolation (RRE) is a classic acceleration method for vector-valued fixed-point process, commonly airsing from iterative solution of algebraic equations. In this talk, we discuss the generalization of this extrapolation framework to sequences of low-rank matrices generated by iterative methods for large-scale matrix equations, such as low-rank alternating directions implicit methods for Lyapunov and Riccati equations, as well as generalized Sylvester equation. We also briefly discuss extrapolation strategies for fixed-point iterations arising in machine learning algorithms.

The Fréchet derivative of the tensor t-function

Authors: Kathryn Lund[1], Marcel Schweitzer[2]

[1] *STFC Scientific Computing*

[2] *University of Wuppertal*

Contact: kathryn.lund@stfc.ac.uk

The tensor t-function, a formalism that generalizes the well-known concept of matrix functions to third-order tensors, is introduced in Lund (Numer Linear Algebra Appl 27(3):e2288). In this work, we investigate properties of the Fréchet derivative of the tensor t-function and derive algorithms for its efficient numerical computation. Applications in condition number estimation and nuclear norm minimization are explored. Numerical experiments implemented by the t-Frechet toolbox hosted at https:// gitlab.com/katlund/t-frechet illustrate properties of the t-function Fréchet derivative, as well as the efficiency and accuracy of the proposed algorithms.

Accurate Models of NVIDIA Tensor Cores

Authors: Faizan Khattak[1], Mantas Mikaitis[1]

[1] *University of Leeds*

Contact: m.mikaitis@leeds.ac.uk

Matrix multiplication is a fundamental operation in both training of neural networks and inference. To accelerate matrix multiplication, Graphical Processing Units (GPUs) provide it implemented in hardware. Due to the increased throughput over the software-based matrix multiplication, the multipliers are increasingly used outside of AI, to accelerate various applications in scientific computing. However, matrix multipliers targeted at AI are at present not compliant with IEEE 754 floating-point arithmetic behaviour, with different vendors offering different numerical features. This leads to non-reproducible results across different generations of GPU architectures, at the matrix multiply-accumulate instruction level. To study numerical characteristics of matrix multipliers - such as rounding behaviour, accumulator width, normalization points, extra carry bits, and others - test vectors are typically constructed. Yet, these vectors may or may not distinguish between different hardware models, and due to limited hardware availability, their reliability across many different platforms remains largely untested.

We present software models for emulating the inner product behavior of low- and mixed-precision matrix multipliers in the V100, A100, H100 and B200 data center GPUs in most supported input formats of interest to mixed-precision algorithm developers: 8-, 16-, and 19-bit floating point. These models enable verification of test vectors before applying them to real hardware and also support computational scientists and mixed-precision algorithm developers with easy-to-use accurate models available in MATLAB - we demonstrate their use in multi-word emulation algorithms for matrix multiplication.

Principal Surjective Flows: Relaxing Bijection Assumption via the Smooth Co-Area Formula and Gram Determinants

Authors: Haoran Ni[1], Martin Lotz

[1] *University of Warwick*

Contact: haoran.ni.1@warwick.ac.uk

Normalizing Flows (NFs) enable tractable density evaluation by modelling data through invertible neural transformations. However, this reliance on global bijectivity severely restricts their expressiveness when the target distribution lies on a low-dimensional manifold or exhibits complex topology. To overcome this limitation, we introduce Principal Surjective Flows (PSFs), a framework that replaces invertibility with carefully constructed surjective mappings. Leveraging the Smooth Co-area Formula, we derive a principled likelihood expression in which the density transformation is governed by the normal Jacobian, which is equivalently the square root of a Gram determinant formed from the Jacobian's nonzero singular values. This reformulation enables stable and geometrically meaningful density computation for dimension-reducing generative maps. Through Gaussian examples, we show that integrating densities over fibres reduces to evaluating Gaussians on affine subspaces determined by the row space and nullspace of the map, revealing a direct connection between surjective flows and core numerical linear algebra concepts including low-rank operators, projection matrices, and conditioning. Overall, PSFs provide a mathematically grounded and computationally efficient extension of normalizing flows, broadening their applicability while preserving exact and tractable density estimation.

Reduced- and Mixed-Precision Algorithms for QR Decomposition

Authors: Eda Oktay[1], Martin Koehler[2], Peter Benner[2], Xiaobo Liu[2]

[1] *Max Planck Institute for Dynamics of Complex Technical Systems Magdeburg*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems*

Contact: oktay@mpi-magdeburg.mpg.de

As machine learning and AI continue to shape modern hardware design, reduced-precision arithmetic has become essential in high-performance computing. Recent advances in hardware architectures—such as AI accelerators, GPUs, and tensor-core technologies—many of which are driven by machine-learning workloads, are optimized for low-precision operations to improve performance and reduce energy consumption in computationally intensive scientific applications. However, the primary drawback of uniformly low-precision methods remains their potential numerical instability, particularly when solving linear systems. As a result, mixed-precision strategies have become widely used in numerical linear and multilinear algebra to balance performance, energy efficiency, and accuracy while still exploiting ML-optimized hardware.

QR factorizations are among the core algorithms in numerical linear algebra and are also playing an increasing role in machine-learning applications such as randomized embeddings, dimensionality

reduction, and least-squares problems. However, their computational cost remains significant for large-scale matrices. In this talk, we first present a uniform half-precision Householder QR decomposition based on the WY representation and examine both its potential performance advantages on modern AI hardware and its significant limitations. We then show how increasing the precision of key computations can substantially improve accuracy and how these choices interact with ML-optimized hardware. Finally, we discuss how reduced- and mixed-precision strategies can be incorporated into other QR decomposition algorithms, including shiftedCholeskyQR3 and TSQR, to further leverage modern AI hardware and software ecosystems.

Fast and explainable clustering in the Manhattan and Tanimoto distance

Authors: Kaustubh Roy[1], Stefan Guettel[1]

[1] *University of Manchester*

Contact: kaustubh.roy@manchester.ac.uk

The CLASSIX algorithm is a fast and explainable approach to data clustering. In its original form, this method utilizes the first principal component of the data matrix to truncate the search for nearby data points, using the Cauchy-Schwarz inequality, with proximity being defined in terms of the Euclidean distance. In this work, we demonstrate methods to extend CLASSIX to other distance measures by showcasing its effectiveness in the Manhattan distance and the Tanimoto distance. CLASSIX in these two distance metrics uses the 1-norm of the data vectors as the sorting criterion. The triangle inequality is used as a general search truncation criterion applicable to any p-norm, and the Baldi intersection inequality is used as a search truncation criterion for the Tanimoto distance.

Numerical linear algebra with neural operator preconditioning for solving some parametric PDEs

Author: Yanfei Xiang[1]

[1] *University of Strasbourg*

Contact: yanfei.xiang@math.unistra.fr

This work exclusively focuses on the mixed precision algorithms that integrate classical numerical linear algebra methods with nonlinear neural network-based preconditioners to accelerate the solution of some parametric Partial Differential Equations (PDEs). Specifically, we consider Krylov subspace methods such as Flexible GMRES (FGMRES) and Flexible FOM (FFOM), combined with nonlinear or flexible preconditioners derived from the trained neural operator models. The parametric PDEs addressed include the Helmholtz equations, Poisson equations, Darcy flow, and

Diffusion-Advection equations, spanning both academic benchmarks and practical dataset. Compared to the classical numerical preconditioners, the trained neural operator-based preconditioners exhibit significant generalization capabilities in addressing a wide range of numerical and parametric variations. These include the ability to address varying source term, speed of sound, diffusivity term, velocity field for advection, and varying boundary conditions for these parametric PDEs. In addition, we support both CPUs and GPUs implementations of the proposed mixed precision algorithms. Overall, this work demonstrates the efficiency and flexibility of combining modern neural network-based solvers with classical Krylov methods, leveraging the strengths of both to achieve higher attainable accuracy and broader applicability across diverse problem settings.

MS 02: New Directions and Challenges in Linear Algebra

Minisymposium Organizer: Ilse Ipsen

The goal is to highlight exciting developments and challenging problems in numerical linear algebra, that have arisen from the influence of theoretical computer science, random matrix theory, high-dimensional probability, and statistical learning theory. Topics include structured matrices, random matrix theory, low-rank approximations, sketching and randomization; and applications to graph theory and scientific machine learning.

List of Talks

Divide-and-Conquer for Nonsymmetric Eigenvalue Problems Part I: Randomization

Authors: Ioana Dumitriu[1], James Demmel[2], Ryan Schneider[2]

[1] *University of California, San Diego*

[2] *University of California Berkeley*

Contact: idumitriu@ucsd.edu

The fastest method for diagonalizing a nonsymmetric matrix or matrix pencil is pseudospectral divide-and-conquer. This two-step algorithm diagonalizes a matrix/pencil by (1) randomly perturbing the input(s) and (2) running fast (and highly-parallel) spectral divide-and-conquer. The key to this approach is the random perturbation, which with high probability implies a guarantee of pseudospectral shattering – i.e., that the spectrum and pseudospectrum of the perturbed problem

is sufficiently well-behaved for divide-and-conquer to succeed and to run efficiently. In this talk, we present the most general formulation of pseudospectral divide-and-conquer and discuss efforts to specialize the algorithm to structured problems (e.g., definite pencils).

Inner product-free approximation of matrix functions

Author: Stefan Güttel[1]

[1] *The University of Manchester*

Contact: stefan.guettel@manchester.ac.uk

Randomized sketching is a promising tool to reduce the number of inner products computed in Krylov methods for solving large systems of linear equations $Ax = b$, or more generally, when approximating the action of a matrix function on a vector, $f(A)b$. For the case of linear systems, it has recently been observed that one can often get away with completely inner product-free Krylov-based preconditioners when a stabilizing outer iteration (like FGMRES) is used. In this talk we discuss whether similar (almost) inner product-free methods can be devised for the $f(A)b$ problem.

Operator learning without the adjoint

Authors: Diana Halikias[1], Alex Townsend[2], Nicolas Boullé[3], Sam Otto[2]

[1] *New York University*

[2] *Cornell University*

[3] *Imperial College, London*

Contact: dh3161@nyu.edu

There is a mystery at the heart of operator learning: how can one recover a non-self-adjoint operator from data without probing the adjoint? Current practical approaches suggest that one can accurately recover an operator while only using data generated by the forward action of the operator without access to the adjoint. However, naively, it seems essential to sample the action of the adjoint. We prove that that without querying the adjoint, one can approximate a family of non-self-adjoint infinite-dimensional compact operators via projection onto a Fourier basis. We then apply the result to recovering Green's functions of elliptic partial differential operators and derive an adjoint-free sample complexity bound.

We also discuss the following question in the discrete, transpose-free setting: for what problems in numerical linear algebra is it necessary to access matrix-vector products with both a matrix and its transpose? This question also arises in the "unmatched backprojector" setting of certain problems

in imaging. In particular, we describe the importance of the transpose in low-rank approximation, least-squares problems, and different varieties of norm estimation.

Column subset selection: A new perspective

Authors: Ilse Ipsen[1], Arvind Saibaba[1]

[1] *North Carolina State University*

Contact: ipsen@ncsu.edu

It is well known that the problem of selecting k columns with maximal volume from a real matrix is NP-hard, and does not admit a polynomial time approximation scheme. However, one can show that the logarithm of the volume is a "submodular" function. Intuitively, this means: Adding a new column to a larger submatrix tends to be less beneficial than adding it to a smaller submatrix.

If subset selection for maximizing the logarithmic volume is implemented via QR with column pivoting, we show that: The Businger Golub algorithm is a greedy algorithm with relative error at most .37, while the Gu-Eisenstat algorithm is a 1-interchange algorithm with relative error at most .5.

Structured Matrix Approximation via Matrix-Vector Products

Authors: Cameron Musco[1], Chinmay Hegde[2], Christopher Musco[3], David Persson[4], Diana Halikias[3], Feyza Duman Keles[3], Noah Amsel[3], Pratyush Avi[2], Tyler Chen[5]

[1] *University of Massachusetts Amherst*

[2] *NYU*

[3] *New York University*

[4] *New York University & Flatiron Institute*

[5] *JP Morgan Chase*

Contact: cmusco@cs.umass.edu

In this talk, I will give an overview of recent progress on the problem of structured matrix approximation from matrix-vector products. Given a target matrix A that can only be accessed through a limited number of (possibly adaptively chosen) matrix-vector products, we seek to find a near-optimal approximation to A from some structured matrix class – e.g., a low-rank approximation, a hierarchical low-rank approximation, a sparse or diagonal approximation, etc. This general problem arises across the computational sciences and data science, both in algorithmic applications and, more recently, in scientific machine learning, where it is closely related to the problem of linear operator learning from input/output samples.

I will overview recent work, where we give 1) optimal algorithms for approximating A with a matrix with a fixed sparsity pattern (e.g., a diagonal or banded matrix), 2) the first algorithms with strong relative error bounds for hierarchical low-rank approximation, and 3) the first bounds for generic structured families with sample complexity depending on the parametric complexity of the family. I will highlight several open questions related to these results.

Divide-and-Conquer for Nonsymmetric Eigenvalue Problems Part II: Implementation

Authors: Ryan Schneider[1], Ioana Dumitriu[2], James Demmel[1]

[1] *University of California Berkeley*

[2] *University of California, San Diego*

Contact: ryan.schneider@berkeley.edu

While pseudospectral divide-and-conquer is optimal for nonsymmetric eigenvalue problems (in terms of both arithmetic and communication complexity) it is not currently implemented in any of our standard numerical linear algebra libraries. This is due to both the difficulty of translating the algorithm's technical pseudocode into something practical and to the challenge of preparing users for a randomized eigensolver, which will necessarily output different eigenvalues each time it runs. This talk explores the obstacles to bringing pseudospectral divide-and-conquer to practice and the reasons for pursuing a high-performance implementation in spite of them.

How ill conditioned can sub-matrices of the Fourier matrix be?

Authors: John Urschel[1], John Tebou[1], Rikhav Shah[1]

[1] *MIT*

Contact: urschel@mit.edu

The discrete Fourier transform matrix, and sub-matrices of it, appear in a wide variety of applications. While the Fourier matrix itself is a scaled unitary matrix, its sub-matrices can be exponentially ill-conditioned. In this talk, we discuss applications, prior work, and we provide tight estimates for just how ill-conditioned such matrices can be.

Spooky Scary Skeletons

Authors: Abinand Gopal[1], Heather Wilber[2], Per-Gunnar Martinsson[3], Wietse Vaes[2]

[1] *UC Davis*

[2] *University of Washington*

[3] *UT Austin*

Contact: hdw27@uw.edu

In low rank approximations to kernel matrices, skeletons consist of subsets of rows and columns from which CUR, ID, and related approximations can be formed. We consider applications in which parameter-dependent families of matrices with numerical low rank structures appear, such as in parameter dependent partial differential equations. We develop new techniques for analyzing and constructing universal skeletons. A universal skeleton is a single skeleton useful for the entire family of matrices; it captures shared low rank structures across the group. We also supply a practical and memory-efficient algorithm for constructing and using universal skeletons via proxy-point methods. To make things sufficiently scary, we develop universal skeletons in the hierarchical setting, and then use them to develop new fast direct solvers for applications involving multifrequency Helmholtz and Lippmann Schwinger equations.

MS 03: Advanced Acceleration and Convergence Techniques for Solving Linear and Nonlinear Systems

Minisymposium Organizer: Yunhui He

This minisymposium showcases recent advances in theoretical analysis and algorithmic development of fast numerical solvers in computational linear algebra. The focus is on scalable methods for large-scale systems, featuring innovations in Krylov subspace techniques, multigrid methods, and acceleration strategies such as Anderson acceleration and nonlinear GMRES. Applications include PDE-constrained optimization, optimal transport, graph inverse problems, multiphysics simulations, and nonlinear polynomial systems, highlighting emerging challenges and opportunities in scientific computing.

List of Talks

Natural Gradient Descent for Hyperparameter Estimation in Bayesian Inverse Problems

Authors: Elle Buser[1], Malena Sabate Landman[2], Julianne Chung[1], Levon Nurbekyan[1]

[1] *Emory University*

[2] *University of Oxford*

Contact: ebuser@emory.edu

We propose a natural gradient descent (NGD) method for efficient hyperparameter estimation in Bayesian inverse problems. In this framework, the objective is formulated as a function of a symmetric positive definite (SPD) matrix that depends on the hyperparameters. Rather than optimizing purely in the parameter space, we define a natural gradient that incorporates information about the model space (i.e. the space of SPD matrices) into the gradient descent by equipping the manifold with a Riemannian metric. We consider two metrics, the Log-Cholesky and Log-Euclidean metric, both designed for SPD matrices. Finally, we present a numerical example comparing the performance of NGD based on different metrics to the standard gradient descent.

A BFBt Preconditioner for Double Saddle-Point Systems

Author: Chen Greif[1]

[1] *The Department of Computer Science*

Contact: greif@cs.ubc.ca

We consider nonsymmetric double saddle-point systems. Given the 3-by-3 block structure of the matrix, the associated block LU decomposition features two Schur complements. A theoretical question we explore is what happens when one of the Schur complements is inverted exactly and the second, nested one, is approximated. Eigenvalue analysis sheds some light on the effect of this type of inexactness on the speed of convergence of minimum residual iterative solvers. Specific analytical and computational observations are made for the case where the nested Schur complement is approximated using the BFBt method.

A Generalized Alternating Anderson Acceleration Method

Authors: Yunhui He[1], Santolo Leveque[2]

[1] *University of Houston*

[2] *Charles University*

Contact: yhe43@central.uh.edu

In this talk, we propose a generalized alternating Anderson acceleration method, a periodic scheme composed of t fixed-point iteration steps, interleaved with s steps of Anderson acceleration with window size m, to solve linear and nonlinear problems. This allows flexibility to use different combinations of fixed-point iteration and Anderson iteration. We present a convergence analysis of the proposed scheme for accelerating the Richardson iteration in the linear case, with a focus on

specific parameter choices of interest. Specifically, we prove convergence of the proposed method under contractive fixed-point iteration and provide a sufficient condition for convergence when the Richardson iteration matrix is diagonalizable and noncontractive. To demonstrate the broader applicability of our proposed method, we use it to accelerate Picard iteration, gradient descent, and the alternating direction method of multipliers in solving partial differential equations and nonlinear, nonsmooth optimization problems. The numerical results illustrate that the proposed scheme is more efficient than the existing windowed Anderson acceleration and alternating Anderson ($s = 1$) in terms of iteration number and CPU time for careful choice of parameters m, s, t.

Accelerated Multigrid Cycles

Authors: Yunhui He[1], Xiaozhe Hu[2], Chunyan Niu[3]

[1] *University of Houston*

[2] *Tufts University*

[3] *Zhengzhou University*

Contact: yhe43@central.uh.edu

Multigrid (MG) methods are efficient and scalable for solving sparse linear systems arising from the discretization of partial differential equations (PDEs). However, the performance of standard V- and W-cycle MG methods often deteriorates as the physical and geometric complexity of the PDEs increases. To remedy this, the Algebraic Multilevel Iteration (AMLI)-cycle was developed, utilizing Chebyshev polynomials to define the coarse-level solver. Despite its theoretical strength, AMLI is not widely used in practice because it requires accurate estimations of extreme eigenvalues and can be difficult to implement.

In this talk, inspired by recent acceleration techniques, we propose new accelerated MG cycles that do not require estimating extreme eigenvalues. We prove that these resulting cycles achieve convergence rates comparable to the Chebyshev-based AMLI-cycle. Furthermore, our approach is more straightforward to implement, making it practical for real-world applications. Numerical experiments are presented to validate our theoretical results.

Robust Numerical Differentiation for Entropy-regularized Optimal Transport (EOT) with application to Shuffled Regression

Authors: Xiaofeng Felix Ye[1], Xingjie Li[2], Yunhui He[3]

[1] *CUNY at albany*

[2] *UNC Charlotte*

[3] *University of Houston*

Contact: xli47@uncc.edu

In this presentation, I will begin by introducing shuffled regression and the entropic optimal transport (EOT) as one possible tool for solving shuffled regression. A common approach for this optimization is to use a first-order optimizer, which requires the gradient of the OT distance. For faster convergence, one might also resort to a second-order optimizer, which additionally requires the Hessian. I will present the analytical derivatives of EOT, provide a brief overview of numerical condition numbers, and explain how to compute a crucial linear system robustly. Through analytical derivation and spectral analysis, we identify the numerical instability caused by the singularity and ill-posedness of a key linear system, prove the asymptotic limits of its condition number when both sample size N goes to infinity and regularization strength ε goes to 0, and improve the efficiency and robustness of computation. Finally, I would like to discuss future work as well as extensions.

GA-NGMRES: An Alternating NGMRES Method for Accelerating First-Order Optimization

Authors: Andreas Mang[1], Yunhui He[2]

[1] *Department of Mathematics, University of Houston*

[2] *University of Houston*

Contact: andreas@math.uh.edu

We propose a generalized alternating nonlinear generalized minimal residual method (GA-NGMRES) for accelerating first-order optimization algorithms. The method is applied to preconditioned first-order schemes by interpreting their update rules as fixed-point iterations. GA-NGMRES introduces a periodic mixing strategy that alternates between NGMRES extrapolation and fixed-point updates, resulting in improved robustness and efficiency. We demonstrate that the proposed approach reduces both iteration counts and overall runtime compared to state-of-the-art methods. Numerical comparisons are provided against preconditioned gradient descent and preconditioned, inexact Gauss–Newton–Krylov methods. Since GA-NGMRES relies only on first-order derivative information, it is straightforward to implement. Performance is evaluated with respect to algorithmic hyperparameters, mesh resolution, and regularization parameters. For the problems considered, GA-NGMRES consistently outperforms Anderson acceleration.

Constant Memory and Synchronization Costs for Nonsymmetric Krylov Methods

Author: Stephen Thomas[1]

[1] *Lehigh University*

Contact: stephethomas@gmail.com

For nonsymmetric operators, GMRES convergence is governed not only by the eigenvalues but also by the field of values and the resolvent norm. Non-normality amplifies the resolvent $||(zI-A)^{-1}||$ far from the spectrum, and the spectral geometry of the numerical range $W(A)$ determines how rapidly GMRES residual polynomials can decrease. The intrinsic information dimension K_∞, a problem-dependent quantity measuring the spectral information a Krylov process can resolve above the floating-point noise floor, provides an effective subspace dimension for non-normal operators.

This talk presents an s-step Newton-Leja GMRES method that addresses non-normality directly. The Newton-Leja polynomial basis generates the Krylov subspace with conditioning controlled by the spectral geometry of $W(A)$, achieving constant memory through the s-step recurrence. NL–GMRES minimises the residual with one global reduction per MGS orthogonalisation step. A polar preconditioner based on the Newton-Schulz iteration applied per block to the basis vectors restores orthogonality without forming the Gram matrix, counteracting the basis degradation caused by non-normality. Together, these components achieve constant memory and bounded synchronisation cost, with convergence governed by the field of values and bounded by K_∞.

Numerical experiments on non-normal convection-diffusion operators illustrate the role of spectral geometry in determining both convergence and basis conditioning.

Algebraic Multigrid for H-Curl-Systems

Authors: Raymond Tuminaro[1], Christian Glusa[1]

[1] *Sandia National Laboratories*

Contact: rstumin@sandia.gov

An algebraic multigrid (AMG) algorithm is proposed for curl-curl electro-magnetics PDEs that are discretized with 1st order edge elements. The key idea behind the algorithm centers on generating edge interpolation operators such that certain near null space properties of the discrete curl-curl operator are preserved on coarse levels so that good AMG convergence rates can be obtained. Specifically, the algorithm guarantees that the algebraically constructed discrete gradient operator lies in the null space of the algebraically generated discrete curl-curl operator on all grid levels of the multigrid hierarchy.

The new algorithm is based on a varient of energy minimization AMG. It takes an already-computed nodal interpolation operator to then define an edge interpolant that satisfies a related discrete commuting relationship while minimizing the energy of the edge interpolation basis functions. Unlike previous works, this algorithm is general in that it can essentially transform any standard nodal interpolation operator into an edge interpolation operator that is suitable for curl-curl problems. Thus, it is possible to leverage traditional AMG schemes to first generate nodal interpolation operators which can then be adapted to edge interpolation operators. We show that the new algorithm can produce "ideal geometric" edge interpolants in some limited cases. Numerical results are pro-

vided showing the overall efficacy, comparing with traditional Reitzinger and Schoberl schemes.

Accelerated local-global coupling for non-isothermal multiphase reactive flow in hydrate-bearing sediments

Authors: Maria Vasilyeva, Ingo Pecher, Richard Coffin, Yusuf Azeez
Contact: maria.vasilyeva@tamucc.edu

We present a dynamic local–global coupling strategy for non-isothermal multiphase reactive flow in hydrate-bearing sediments, where hydrate phase change strongly couples transport, pressure, and temperature through variations in porosity and permeability. The nonlinear system is solved by sequential Picard iterations with physics-based splitting into transport, flow, and heat processes. To reduce the cost of global iterations, we introduce dynamic nonlinearity localization. Reaction-based markers identify active subdomains, which are enlarged by oversampling (buffer cells). A locally accelerated nonlinear solve is performed in these regions before a global correction step.

---

nlKrylov: A unified framework for nonlinear GCR-type Krylov subspace methods

Authors: Agnieszka Miedlar[1], Ning Wan[1], Tom Werner[2]

[1] *Virginia Tech*

[2] *TU Braunschweig*

Contact: tom.werner@tu-braunschweig.de

In this talk, we introduce a unified framework for nonlinear Krylov subspace methods (nlKrylov) to solve systems of nonlinear equations. Building on the recent development of nlTGCR as well as earlier work on classical GCR-like linear Krylov solvers such as GMRESR, we generalize these approaches to non-linear problems via nested algorithmic structures. We establish connections of nlKrylov methods to other existing nonlinear methods such as quasi-Newton and subspace projection methods. Our theory is completed by rigorous convergence results for problems with both nonsingular and singular Jacobian. The framework is further extended to matrix-valued rootfinding problems using global nonlinear Krylov approaches. Extensive numerical experiments validate the theoretical insights and demonstrate the robustness and efficiency of our proposed algorithms.

Flattening and Middle Rank Tensor Approximation

Authors: Zequn Zheng[1], Hongchao Zhang[1]

[1] *Louisiana State University*

Contact: zzheng@lsu.edu

Canonical polyadic tensor decomposition and approximation are fundamental problems in multilinear algebra with broad applications in signal processing, machine learning, and scientific computing. The difficulty of those problems depends on both the rank and order of the tensors. We introduce a new method for middle-rank tensor approximation and prove a new criterion for reshaping higher-order tensors. Our method leverages generating polynomials and utilizes linear algebra to generate a good starting point for the middle rank tensor approximation problem. When the given tensor is sufficiently close to a tensor whose rank is below a certain bound, we prove that our algorithm gives a quasi-optimal tensor approximation. Numerical experiments demonstrate that our algorithm can produce accurate tensor approximations for order 3 and higher-order tensors.

MS 04: Matrix Geometries

Minisymposium Organizer: Rongbiao Wang

The connection between geometry and linear algebra is a profound subject with numerous real-world applications spanning machine learning, statistics, robotics, and optimization. This minisymposium brings together researchers who leverage geometric structure to explore the two-way interplay between geometry and matrices. On one hand, matrices are used to construct new models for geometric quantities that are more easily handled and computationally tractable. At the same time, open problems in areas such as applied algebraic geometry have been addressed through the lens of matrix theory. On the other hand, geometry provides new insights into matrix theory. This has long been a central theme, from the study of eigenvalue crossings via characteristic classes to recent developments in matrix decompositions and algorithms inspired by geometric tools, including factorizations of matrices into products of manifold-represented factors. Moreover, advances in random matrix theory continue to reveal deep connections to differential geometry. With this variety of expertise, the minisymposium aims to investigate how geometry informs all stages of research in matrix theory and to surface unifying principles that bridge these perspectives.

List of Talks

Locally Diffeomorphic Logarithm of Special Orthogonal Matrices

Authors: Zhifeng Deng[1], Kyle Gallivan[2], Pierre-Antoine Absil[3], Wen Huang[1]

[1] *Xiamen University*

[2] *Florida State University*

[3] *UCLouvain*

Contact: zhifengdeng@xmu.edu.cn

The special orthogonal group $\mathbb{SO}_n$ is a Lie group whose geometry and local structure are encoded by the exponential map on its Lie algebra $\mathbf{Skew}_n$, the set of skew-symmetric matrices. The associated inverse problem—the matrix logarithm—exhibits a highly nontrivial local diffeomorphism structure, and the notion of a nearby logarithm arises naturally as a local inverse of the skew-restricted exponential. In this work, we establish a complete description of the local diffeomorphism structure of the exponential on the complement of tangent conjugate locus, the skew-symmetric matrices where the derivative of the exponential is rank-deficient. We show that the connected components of this complement admit a systematic organization into countably many diffeomorphic regions. By introducing a canonical alignment rule for Schur decompositions, we obtain a labeling and interpretation of these components. Based on this geometric framework, we propose an efficient and numerically stable algorithm for computing the nearby logarithm with two Schur decompositions. We further show that whenever a nearby logarithm exists, it coincides with the nearest Frobenius-norm preimage, providing a practical and reliable alternative for local inversion.

On the Identifiability of Transformer Self-Attention

Author: Nathan Henry[1]

[1] *University of California, Berkeley*

Contact: nathan.henry@berkeley.edu

Multi-head self-attention is a fundamental building block of the transformer architecture in modern machine learning, enabling large language models and much of modern generative AI as we know it. However, some aspects of the self-attention function space remain poorly understood. In particular, its parameterization is non-unique: continuous families of unique weight matrices can induce the same input–output map. This talk studies the self-attention function space, analyzing this non-uniqueness.

We study two common variants of self-attention. In linear attention, the attention map is a matrix of cubic polynomials in the inputs and tri-linear in the parameters, making it natural to use tools from algebraic geometry to describe generic fibers in the space. In softmax attention, a normalization breaks symmetries, leading to a different notion of generic identifiability, which we study with complex analysis.

Special orthogonal, special unitary, and symplectic groups as products of Grassmannians

Authors: Lek-Heng Lim[1], Ke Ye[2], Xiang Lu[3]

[1] *The University of Chicago*

[2] *the Academy of Mathematics and Systems Science, Chinese Academy of Sciences*

[3] *University of Chicago*

Contact: xiangl1@uchicago.edu

We describe a curious structure of the special orthogonal, special unitary, and symplectic groups that has not been observed, namely, they can be expressed as matrix products of their corresponding Grassmannians realized as involution matrices. We will show that $\mathrm{SO}(n)$ is a product of two real Grassmannians, $\mathrm{SU}(n)$ a product of four complex Grassmannians, and $\mathrm{Sp}(2n, \mathbb{R})$ or $\mathrm{Sp}(2n, \mathbb{C})$ a product of four symplectic Grassmannians over $\mathbb{R}$ or $\mathbb{C}$ respectively.

Stochastic approximations with operator means

Author: Pálfia Miklós[1]

[1] *Corvinus University of Budapest*

Contact: miklos.palfia@uni-corvinus.hu

In this talk we investigate zeros of nonlinear operators on the cone of positive definite operators over a Hilbert space. The unique zero of such nonlinear operators define means of positive definite operators. Moreover these nonlinear operators generate strictly exponentially contracting semigroups in some metric defined on positive operators. We survey recent results that establish the operator norm convergence of deterministic and stochastic resolvent and proximal type algorithms, in particular versions coming from a Trotter-Kato type splitting formula. Applications include generalizations of results proved under strong moment conditions for generalized Karcher means. The talk is based on recent joint work with Zoltán Léka and is a continuation of our earlier project.

Differential-Geometric View of the Schur–Horn Theorem and Related Convexity Phenomena

Author: Tin-Yau Tam[1]

[1] *University of Nevada, Reno*

Contact: ttam@unr.edu

We present a differential–geometric view of the Schur–Horn theorem and related convexity phenomena. For an $n \times n$ Hermitian matrix A with simple spectrum, the Schur–Horn map

$$\mu : \mathrm{U}(n) \to \mathbb{R}^n, \quad \mu(U) = \mathrm{diag}(UAU^{-1}), \qquad U \in \mathrm{U}(n),$$

is shown to be a proper submersion over the relative interior of the Schur–Horn polytope, where $\mathrm{U}(n)$ is the unitary group. We obtain a smooth path-lifting property and a global smooth selection along

any line segment in the interior, providing a geometric strengthening of the classical majorization theorem and a proof of Westwick's c-numerical range theorem. The talk will emphasize the matrix geometry and topology.

Stiefel Optimization is NP-hard

Authors: Tianyun Tang[1], Lek-Heng Lim[1], Zehua Lai[2]

[1] *The University of Chicago*

[2] *University of Texas at Austin*

Contact: ttydp123@gmail.com

We show that linearly constrained linear optimization over a Stiefel or Grassmann manifold is NP-hard in general. We show that the same is true for unconstrained quadratic opti- mization over a Stiefel manifold. We will show that unless P = NP, these optimization problems over a Stiefel manifold do not have FPTAS. As an aside we extend our results to flag manifolds. Combined with earlier findings, this shows that manifold optimization is a difficult endeavor — even the simplest problems like LP and unconstrained QP are already NP-hard on the most common manifolds.

MS 05: Matrix Nearness Problems

Minisymposium Organizers: Volker Mehrmann, Vanni Noferini

The characterization and numerical solution of matrix nearness problems, such as the distance to instability, non-passivity, or singularity, constitute a highly active research area. In recent years, new approaches based on gradient flows and Riemannian optimization have emerged, offering promising alternatives to traditional methods based on eigenvalues and singular values. These developments are particularly important for large scale problems, where such techniques must often be combined with model reduction or Krylov subspace methods. This minisymposium will present the state of the art in this field, along with recent advances in computational methods. In addition to results for classical unstructured matrices and matrix pairs, structured cases will also be discussed.

List of Talks

Distance to prescribed rank matrix polynomials via generic factorizations and alternating least squares

Authors: Andrii Dmytryshyn[1], Froilán M. Dopico[2], Marten Gulliksson[3], Rakel Hellberg[4], Sweta Das[3]

[1] *Chalmers University of Technology and University of Gothenburg (Sweden)*

[2] *Universidad Carlos III de Madrid (Spain)*

[3] *Örebro University (Sweden)*

[4] *Chalmers University of Technology and University of Gothenburg (Sweden)*

Contact: dopico@math.uc3m.es

We propose an algorithm that approximates a given matrix polynomial of any degree d by another matrix polynomial of a prescribed rank and degree at most d. The algorithm combines recent advances in the theory of generic factorizations for matrix polynomials of bounded rank and degree with an alternating least squares strategy. For $d = 1$, the algorithm includes the important case of matrix pencils. The algorithm solves, as a particular case, the well-known problem of computing the distance to singularity of a regular matrix polynomial. We present numerical experiments for testing the proposed algorithm and for comparison to the previously known ones.

A Newton–bisection method with monotone convergence for matrix nearness problems

Authors: Nicola Guglielmi[1], Christian Lubich[2]

[1] *Gran Sasso Science Institute*

[2] *Universitaet Tuebingen*

Contact: nicola.guglielmi@gssi.it

In this talk, I will address eigenvalue-optimization-based matrix nearness problems such as the stability radius of a matrix or a time invariant system, $\mathcal{H}^\infty$ norm computation, the structured distance to singularity.

These are formulated here as 2-variable optimization problems of functionals depending either on a single or on several target eigenvalues of the matrix. It is well-known that the classical Newton-bisection method (proposed, e.g., in [1]) may fail to converge.

For this, in 2016, Mitchell and Overton [3] proposed a very effective convergent alternating iteration method, the Hybrid Expansion Contraction (HEC), which was then extended to passivity optimization in [4].

The method we propose here uses a two-level nested iteration, where, at the inner level, a constrained gradient system is integrated to optimize the functional over the set of matrices of specified norm and, at the outer level, to compute the optimal norm. However, an important modification of the method exploits the free gradient system, allowing for guaranteed convergence (see [2]). A discussion and a comparison wrt the HEC methodology will conclude the talk.

References.
[1] N. Guglielmi, M. Gürbüzbalaban, and M.L. Overton. Fast approximation of the $\mathcal{H}^\infty$ norm via optimization over spectral value sets. SIAM J. Matrix Anal. Appl., 34(2): 709–737, 2013.

[2] N. Guglielmi and C. Lubich. Matrix nearness problems and eigenvalue optimization.

[3] T. Mitchell and M. L. Overton. Hybrid expansion–contraction: a robust scalable method for approximating the $\mathcal{H}^\infty$ norm. IMA J. Numer. Anal., 36(3): 985–1014, 2016.

[4] T. Mitchell and P. Van Dooren. Root-max problems, hybrid expansion-contraction, and quadratically convergent optimization of passive systems. SIAM J. Matrix Anal. Appl., 44(2): 753–780, 2023.

Bounds on the geodesic distances on the Stiefel manifold for a family of Riemannian metrics

Authors: Simon Mataigne[1], Pierre-Antoine Absil[1], Nina Miolane[2]

[1] *UCLouvain*

[2] *UCSB*

Contact: simon.mataigne@uclouvain.be

We give bounds on geodesic distances on the Stiefel manifold, derived from new geometric insights. The considered geodesic distances are induced by the one-parameter family of Riemannian metrics introduced by Hüper et al. (2021), which contains the well-known Euclidean and canonical metrics. First, we give the best Lipschitz constants between the distances induced by any two members of the family of metrics. Then, we give a lower and an upper bound on the geodesic distance by the easily computable Frobenius distance. We give explicit families of pairs of matrices that depend on the parameter of the metric and the dimensions of the manifold, where the lower and the upper bound are attained. These bounds aim at improving the theoretical guarantees and performance of minimal geodesic computation algorithms by reducing the initial velocity search space. In addition, these findings contribute to advancing the understanding of geodesic distances on the Stiefel manifold and their applications.

Robustly asymptotically stable dissipative Hamiltonian descriptor systems

Authors: Volker Mehrmann[1], Peter Benner[2], Anshul Prajapati[3], Punit Sharma[4]

[1] *TU Berlin*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems*

[3] *MPI for Dynamics of Complex Technical Systems Magdeburg*

[4] *Indian Institute of Technology Delhi*

Contact: mehrmann@math.tu-berlin.de

We analyze the robust asymptotic stability under structure-preserving perturbations for the class of linear time-invariant dissipative-Hamiltonian differential-algebraic (dHDAE) systems. We show how to compute the distance to the nearest singular and high index system and determine stability radii for the finite spectrum under structure preserving perturbations.

Singular Value Characterizations for a Nearest Rectangular Polynomial Matrix with an Eigenvalue

Authors: Emre Mengi[1], Lauri Nyman[2], Paul Van Dooren[3], Tolga Ucar[1], Vanni Noferini[4]

[1] *Koc University, Istanbul*

[2] *University of Manchester*

[3] *Universite catholique de Louvain*

[4] *Aalto University*

Contact: emengi@ku.edu.tr

A rectangular polynomial matrix with more columns than rows generically has no eigenvalues. We aim to find a smallest perturbation (with respect to the 2-norm of the concatenated coefficients of the polynomial matrix) so that the perturbed polynomial matrix has an eigenvalue, that is prescribed. This problem is motivated by locating a nearest uncontrollable system for a first-order, as well as a higher-order continuous-time control system. We consider the cases when the polynomial matrix is subject to complex and real perturbations. For both, we derive singular value optimization characterizations that facilitate locating a nearest polynomial matrix with the prescribed eigenvalue. The real perturbation case, inspired by the work of Qiu et al. [Automatica, Vol. 31, pp. 879-890] for the distance to instability of a linear continuous-time system under real perturbations, is much more involved, and leads to a more complicated singular value formula.

We exploit the derived singular value optimization characterizations to compute a nearest rectangular polynomial matrix with an eigenvalue under complex perturbations and under real perturbations. The approaches that we devise are based on level-set methods that date back to Byers, Boyd-Balakrishnan, Bruinsma-Steinbuch for the distance to instability, as well as Lipschitz-continuity based global optimization techniques. They seem to work extremely effectively. For instance, we are able to compute a nearest uncontrollable system for a first-order or higher-order system of medium size in a few seconds.

Reliably Computing the Worst-case H-infinity Norm of a Parametric System Using an Interpolation-based Algorithm

Authors: Tim Mitchell[1], Peter Benner[2]

[1] *CUNY Queens College & The Graduate Center*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems*

Contact: tmitchell@qc.cuny.edu

In this talk, we consider computing the worst-case (highest) $\mathcal{H}_\infty$ norm of a either a continuous-time or discrete-time linear time-invariant parametric system, where the state-space matrices all depend on a single real-valued scalar parameter p on a domain $\mathcal{P}$ consisting of a finite number of intervals. On each interval in $\mathcal{P}$, we assume that p may vary nonlinearly and that each state-space matrix is differentiable with respect to p almost everywhere, but we do permit that the parameter may be discontinuous at a finite number of points inside each interval of $\mathcal{P}$. Using the framework of interpolation-based globality certificates, which were first devised as a fast way to compute Kreiss constants and the distance to uncontrollability, we present a new optimization-with-restarts algorithm that, under reasonable assumptions, solves the underlying nonconvex optimization with global convergence and thus computes the worst-case $\mathcal{H}_\infty$ norm to arbitrary accuracy. Our algorithm does $\mathcal{O}(kn^3)$ work, where n is the dimension of the state vector and k is the total number of function evaluations needed to globally approximate our certificate functions over the parameter domain $\mathcal{P}$ using, say, Chebfun. In practice, k is not strongly correlated with n and typically only a handful of optimization restarts are needed before a global optimizer attaining the value of the worst-case $\mathcal{H}_\infty$ norm is found. Moreover, the overall cost of our algorithm is almost entirely due to the cost of approximating the final certificate function, which asserts that the final maximizer found is in indeed a global maximizer. Experiments show that our new algorithm is both significantly faster and more reliable than other reasonable approaches for computing the worst-case $\mathcal{H}_\infty$ norm.

Nearest matrix with multiple eigenvalues by Riemannian optimization

Authors: Vanni Noferini[1], Federico Poloni[2], Lauri Nyman[3]

[1] *Aalto University*

[2] *Università di Pisa*

[3] *University of Manchester*

Contact: vanni.noferini@aalto.fi

Given a square complex matrix A, we tackle the problem of finding the nearest matrix with multiple eigenvalues or, equivalently when A had distinct eigenvalues, the nearest defective matrix. To this goal, we extend the general framework described in [M. Gnazzo, V. Noferini, L. Nyman, F. Poloni, Riemann-Oracle: A general-purpose Riemannian optimizer to solve nearness problems in matrix theory, Found. Comput. Math. 2025] and based on variable projection and Riemannian

optimization, allowing the ambient manifold to simultaneously track left and right eigenvectors. Our method also allows us to impose arbitrary complex-linear constraints on either the perturbation or the perturbed matrix; this can be useful to study structured eigenvalue condition numbers. We present numerical experiments, comparing with preexisting algorithms.

Structured stability radii of dissipative Hamiltonian systems

Authors: Anshul Prajapati[1], Peter Benner[1], Punit Sharma[2], Volker Mehrmann[3]

[1] *Max Planck Institute for Dynamics of Complex Technical Systems*

[2] *Indian Institute of Technology Delhi*

[3] *TU Berlin*

Contact: prajapati@mpi-magdeburg.mpg.de

We study linear time-invariant Dissipative Hamiltonian (DH) systems arising in energy-based modeling of dynamical systems. An advantage of DH systems is that they are always stable due to the structure of their coefficient matrices, and, under further weak conditions, even asymptotically stable. Here, we consider the computation of the stability radii for a given asymptotically stable DH system; i.e., the smallest structured perturbation that puts a DH system on the boundary of the region of asymptotic stability, so that it has purely imaginary eigenvalues. We obtain explicit computable formulas for various structured stability radii. For this, the problem of computing stability radii is reformulated in terms of minimizing the Rayleigh quotient of a Hermitian matrix or the sum of two generalized Rayleigh quotients of Hermitian semidefinite matrices. This reformulation results in the problem of minimizing the largest eigenvalue of an eigenvector-dependent Hermitian matrix or minimizing the smallest eigenvalue of a Hermitian matrix which depends on the eigenvector. We then, demonstrate (via numerical experiments) that, under structure-preserving perturbations, the asymptotic stability of a DH system is much more robust than under general perturbations, since the distance to instability is typically much larger when structure-preserving perturbations are considered.

Constrained Rayleigh quotient optimization and its applications in polynomial eigenvalue problem

Authors: Punit Sharma[1], Anshul Prajapati[2], Ritwik Prabin Kalita[1]

[1] *Indian Institute of Technology Delhi - Abu Dhabi, UAE*

[2] *MPI for Dynamics of Complex Technical Systems Magdeburg*

Contact: punit1718@iitdabudhabi.ac.ae

Let $G, H_1, \ldots, H_s \in \mathbb{C}^{n,n}$ be Hermitian and $S_1, \ldots, S_k \in \mathbb{C}^{n,n}$ be symmetric matrices. In this talk, we maximize the Rayleigh quotient of a Hermitian matrix H under certain constraints involving Hermitian matrices $H_1, \ldots, H_s$ and symmetric matrices $S_1, \ldots, S_k$. More specifically, we compute

$$\mu(G, H_1, \ldots, H_s, S_1, \ldots, S_k) := \sup \Big\{ \tfrac{v^*Gv}{v^*v} : \ v \in \mathbb{C}^n \setminus \{0\},\ v^*H_iv = 0, \\ i = 0, \ldots, s,\ v^TS_jv = 0,\ j = 1, \ldots, k \Big\},$$

where T and $*$ denote the transpose and the conjugate transpose of a matrix or a vector, respectively. Such problems occur in stability analysis of uncertain systems and in the eigenvalue perturbation theory of matrices and matrix polynomials. The problem without symmetric conditions has been studied by [Bora et al., SIMAX, 35(2), 2014] and a computable formula for $\mu(G, H_1, \ldots, H_s)$ was obtained in terms of the largest eigenvalue of a parameter-dependent Hermitian matrix. Similarly, this problem without Hermitian conditions was considered by [Prajapati & Sharma, LAA, 645, 2022] and a computable formula $\mu(G, S_1, \ldots, S_k)$ was obtained in terms of the second largest eigenvalue of some parameter-dependent Hermitian matrix.

In this talk, we extend the ideas from [Bora et al., SIMAX, 35(2), 2014] and [Prajapati & Sharma, LAA, 645, 2022] and derive a computable estimation for $\mu(G, H_1, \ldots, H_s, S_1, \ldots, S_k)$ with both Hermitian and symmetric constraints. This estimation is exact when the eigenvalue at the optimal is simple. We then apply these results in computing the structured eigenvalue backward errors of real T-even and real T-odd matrix polynomials. These results are also shown to be useful in obtaining the eigenvalue backward errors of matrix functions involving both Hermitian and skew-symmetric matrices.

MS 06: Matrix Inequalities, Matrix Equations, and Their Applications

Minisymposium Organizers: Tin Yau Tam, Dominique Guillot, Mohsen Aliabadi

The organizers propose a minisymposium titled "Matrix Inequalities, Matrix Equations, and Their Applications" at ILAS 2026. The goal of this session is to stimulate research and foster interaction among researchers working in these vibrant and rapidly evolving areas. In recent years, research on matrix inequalities, matrix equations, and their applications has achieved significant progress, deepening our theoretical understanding and expanding practical uses across diverse disciplines. This minisymposium will provide a collaborative platform for researchers to present their latest results, exchange innovative ideas, and discuss emerging trends and applications. By bringing together established experts and early-career researchers, the session aims to promote dynamic discussions, inspire new directions, and encourage fruitful collaborations that advance these important fields of study. We have 18 speakers from 6 countries, ranging from PhD students and lecturers to postdocs and senior professors.

List of Talks

Stabilizer Fields and Dimension Growth in Product-Spans

Author: Mohsen Aliabadi[1]

[1] *Clayton State University*

Contact: mohsenmath88@gmail.com

We survey classical results in additive combinatorics and develop linear analogues over field extensions, with an emphasis on Kneser-type phenomena. In addition to recalling Kneser's theorem and stabilizer methods (including Cauchy–Davenport and DeVos's refinement), we present a rigidity theorem showing that if $|A+B| = |A| + |B| - 1$ with $A + B \neq G$, then $A + B$ is a subgroup and A is a coset; an n-fold Kneser bound expressed in terms of stabilizers; a coset-sparsity condition that yields near Cauchy–Davenport bounds; and a density estimate guaranteeing large sumsets.

We then establish linear analogues for finite-dimensional K-subspaces $A, B \subseteq L$: the K-span $\langle AB \rangle$ admits stabilizers that are intermediate fields, leading to a linear rigidity theorem and sharpened Hou–Leung–Xiang bounds under separability.

We conclude with open problems on weakening separability assumptions and on bridging the group and linear settings via stabilizer geometry and coset projections.

Pressing sequences in nonbinary fields

Authors: Declan Liming[1], Hays Whitlatch[2], Joshua Cooper[1]

[1] *University of South Carolina*

[2] *Gonzaga University*

Contact: cooper@math.sc.edu

Pressing sequences of simple undirected graphs totally colored by a field arise naturally in computational phylogenetics, where (over $GF(2)$) they are in bijection with sortings-by-reversal of signed permutations which model gene sequences in related organisms. They can be viewed in several surprising different ways, including as any consecutive initial sequence of rows whose diagonal elements can serve, in natural order, as the pivots of Gaussian elimination performed on a symmetric matrix over the chosen field, with the caveat that every pivot entry is a nonzero square. This gives rise to their connection with positivity: for example, a symmetric matrix M is positive definite iff the graph of which it is the (colored) adjacency matrix is "pressable" in the order given by the rows of M. Thus, finding such an ordered basis of the row space with this property is a natural question, and one for which we have a simple polynomial time algorithm when the base field is $GF(2)$. However, the only algorithms known for fields of characteristic greater than 2 have time complexity $O(n!)$ where n is the number of vertices in the graph (and the dimension of the corresponding matrix). We discuss progress on this question and future directions.

T-Eigenvalues of Third-order Quaternion Tensors

Author: Meiling Deng[1]

[1] *University of Nevada, Reno*

Contact: meilingd@unr.edu

In this talk, we investigate the eigenvalue problem for third-order quaternionic tensors. We first introduce the notion of right T-eigenvalues and develop an efficient algorithm for their computation, whose effectiveness is demonstrated through comparative numerical experiments.

For Hermitian quaternionic tensors, we then derive bounds for the eigenvalues of tensor sums and extend Weyl's classical theorem from matrices to the tensor setting.

Finally, we generalize the Geršgorin disc theorem to obtain eigenvalue localization results for such tensors, providing a practical tool for spectral estimation.

Sharp lower bounds for generalized operator products

Authors: Dominique Guillot[1], Javad Mashreghi[2], Prateek Kumar Vishwakarma[3]

[1] *University of Delaware*

[2] *Université Laval*

[3] *Universite Laval, Quebec, Canada*

Contact: dguillot@udel.edu

We consider general bilinear products parameterized by positive semidefinite matrices. Typically non-commutative, non-associative, and non-unital, these products preserve positivity and include the classical Hadamard, Kronecker, and convolutional products as special cases. We prove that every such product satisfies a sharp nonzero lower bound in the Loewner order, generalizing previous results of Vybíral and Khare that were obtained in the special case of the Hadamard product. Our results naturally extend to Hilbert spaces for a family of products parametrized by positive trace-class operators, providing a lower bound in the Loewner order for such general products, including for the Hilbert tensor product.

(Joint work with Javad Mashreghi and Prateek Kumar Vishwakarma.)

Inequalities for different means of positive definite matrices

Authors: Tanvi Jain[1], Rajendra Bhatia, Yongdo Lim

[1] *Indian Statistical Institute*

Contact: tanvi@isid.ac.in

We discuss majorisation inequalities for different means of positive definite matrices focusing on the geometric mean, the Wasserstein mean, the log Euclidean mean and the power mean.

Riccati Differential Equations, State Transition Matrices, and State Covariance Matrices

Authors: Fengjiao Liu[1], Panagiotis Tsiotras[2], Xiaoyang Ming[1], Yixiao Zhang[1]

[1] *FAMU-FSU College of Engineering*

[2] *Georgia Institute of Technology*

Contact: fliu@eng.famu.fsu.edu

In this talk, we first investigate the maximal interval of existence for the solution of a symmetric matrix Riccati differential equation. Then, we apply this result to study the reachability of the closed-loop state transition matrix for a linear time-varying system over a finite time interval. Under a mild assumption, we characterize the set of closed-loop terminal state transition matrices reachable from the identity matrix using controls of the state feedback form. Lastly, we provide the set of terminal state covariance matrices reachable from a given positive definite initial state covariance matrix when the linear system is not necessarily controllable.

Random Matrix Ensembles with Split Limiting Behavior

Author: Steven Miller[1]

[1] *Williams College*

Contact: sjm1@williams.edu

Random matrix theory has successfully modeled a variety of systems, from energy levels of heavy nuclei to zeros of the Riemann zeta function. One of the central results is Wigner's semi-circle law: the distribution of normalized eigenvalues for ensembles of real symmetric matrices converge to the semi-circle density (in some sense) as the matrix size tends to infinity. We introduce a new family of $N \times N$ random real symmetric matrix ensembles, the k-checkerboard matrices, whose limiting spectral measure has two components. All but k eigenvalues are in the bulk, and their behavior, appropriately normalized, converges to the semi-circle as $N \to \infty$; the remaining k are tightly constrained near N/k and their distribution converges to the $k \times k$ hollow GOE ensemble (this is the density arising by modifying the GOE ensemble by forcing all entries on the main diagonal to be zero). Similar results hold for complex and quaternionic analogues. We are able to isolate each regime separately through appropriate choices of weight functions for theeigenvalues and then an analysis of the resulting combinatorics. This is joint work with Paula Burkhardt, Peter Cohen,

Jonathan Dewitt, Max Hlavacek, Carsten Sprunger, Yen Nhi Truong Vu, Roger Van Peski, and Kevin Yang.

On a characterization of the equality case in the generalized Böttcher-Wenzel inequality

Author: Motoyuki Nobori[1]

[1] *Graduate School of Science and Engineering, Ehime University*

Contact: m819002c@mails.cc.ehime-u.ac.jp

The main purpose of this presentation is to illustrate an equivalent condition for the equality case in the generalized Böttcher-Wenzel (BW) inequality for three matrices. Let A, B, and C be square matrices with complex elements. The BW inequality is an upper bound estimate on the Frobenius norm of the commutator of A and B, defined as $AB - BA$. After the BW inequality was proved in 2008, various subsequent problems related to it have been considered until now. An necessary and sufficient condition for the equality case in the inequality is one of them. Also, there are some papers which investigate relationships between the BW inequality and the uncertainty principle in quantum physics. Recently, the presenter estimated the Frobenius norm of the generalized commutator of A, B, and C, defined as $ABC - CBA$, and obtained an inequality, which is a generalization of the BW inequality [M. Nobori, A generalization of the Böttcher-Wenzel inequality for three rectangular matrices, Linear Algebra Appl. 725 (2025) 135–144]. In this talk, we shall deduce an equivalent condition for the equality case in the generalized BW inequality. The core idea of the derivation is a necessary and sufficient condition for the Weyl inequality.

From binary symmetric matrices to Coxeter-like graphs and self-dual codes

Authors: Marko Orel[1], Draženka Višnjić

[1] *University of Primorska*

Contact: marko.orel@upr.si

The matrix equation $rank(A - B) = 1$ is well studied in linear algebra and combinatorics within *preserver problems* and the theory of *distance-regular graphs/association schemes*. In the talk I will present how this equality is related to coding theory, namely to binary *self-dual codes*.

Let $\widehat{\Gamma}_n$ be the graph with the vertex set formed by all $n \times n$ symmetric matrices with coefficients in the binary field $\mathbb{F}_2 = \{0, 1\}$ where two matrices A and B form an edge if and only if $rank(A-B) = 1$. Graph $\widehat{\Gamma}_n$ was studied in numerous papers and the distance function $d_{\widehat{\Gamma}_n}(A, B)$ between arbitrary vertices is well known and easy to compute. On the other hand, its subgraph Γ_n, which is induced by

invertible matrices, has not attained much attention till recently. In fact, graph Γ_n was introduced in 2015 [1]. It generalizes the well-known Coxeter graph (obtained if $n = 3$).

Recently, we computed the distance function $d_{\Gamma_n}(A, B)$ [2], which turns out to be related to coding theory. Namely, for odd $n \geq 3$, each linear self-dual code C in $\mathbb{F}_2^{n+1}$ can be identified with a certain subset $\mathcal{F}_C$ of the set $\mathcal{SD}_n = \left\{A \in V(\Gamma_n) : d_{\Gamma_n}(A, I) = \frac{n+5}{2}, rank(A - I) = \frac{n+1}{2}\right\}$ where I is the identity matrix. In particular, the matrices $A \in V(\Gamma_n)$, which represent self-dual codes, are fully determined by the values of two graph parameters: $d_{\Gamma_n}(A, I)$ and $d_{\widehat{\Gamma}_n}(A, I)$. In the talk, I will describe the identification $C \leftrightarrow \mathcal{F}_C$ in more details.

References

[1] M. Orel, On generalizations of the Petersen and the Coxeter graph. Electron. J. Combin. 22(4) (2015), Paper #P.4.27.
[2] M. Orel, D. Višnjić, The distance function on Coxeter-like graphs and self-dual codes. Finite Fields Appl. 103 (2025), paper 102580, 51 pp., https://doi.org/10.1016/j.ffa.2025.102580

Permutations induced by GEPP

Author: John Peca-Medlin[1]

[1] *UCSD*

Contact: jpecamedlin@ucsd.edu

Gaussian elimination with partial pivoting (GEPP) remains the most widely used dense linear solver. GEPP produces the factorization $PA = LU$, where L and U are lower and upper triangular matrices and P is a permutation matrix; together, these encode the pivoting strategy, directly influencing stability through classical growth-factor bounds and matrix norm inequalities. When A is random, the permutation arising from the P factor is itself random. When is this permutation uniform? How many cycles appear in its disjoint cycle decomposition (equivalently, how many pivot movements occur during GEPP — providing a new link between random matrix theory and permutation statistics through numerical linear algebra)? What is the length of the longest increasing subsequence of this permutation? We provide statistical answers to these questions for selected random matrix ensembles and transformations. For structured cases involving butterfly permutations, we present full distributional descriptions of these statistics. Moreover, we introduce a random butterfly matrix ensemble whose GEPP-induced permutations are distributed according to Haar measure on a full 2-Sylow subgroup of the symmetric group on a set of size 2^n.

Revisiting the Upper Bild Convexity of Quaternionic Numerical Range

Authors: Jie Tian[1], Tin-Yau Tam[1]

[1] *University of Nevada, Reno*

Contact: jtian@unr.edu

The quaternionic numerical range of a matrix is generally nonconvex, in contrast to the classical complex case. Nevertheless, a theorem of So and Thompson in 1996 asserts that the associated *upper bild* in the complex upper half-plane is always convex.

The original proof of So and Thompson relies on a detailed case-by case and computationally involved analysis, including a reduction to the 2×2 case and an explicit description of boundary curves. A recent note by Kumar in 2019 proposes a shorter and more conceptual argument.

In this talk, we revisit Kumar's approach and identify several gaps in the argument, showing that certain steps require additional justification. These issues point to the importance of incorporating geometric structure into the analysis of the upper bild. We also discuss ongoing work exploring algebraic and geometric features of the upper bild toward a more conceptual understanding of its convexity.

[1] W. So and R. C. Thompson, Convexity of the upper complex plane part of the numerical range of a quaternionic matrix, Linear Multilinear Algebra 41 (1996) 303–362.

[2] P. S. Kumar, A note on convexity of sections of quaternionic numerical range, Linear Algebra Appl. 572 (2019) 92–116.

Convolution of matrices: Cayley–Hamilton theory, matrix transforms, and positivity preservers, with connections to the Bruhat order

Authors: Javad Mashreghi[1], Mostafa Nasri[2], Prateek Kumar Vishwakarma[3]

[1] *Universite Laval*

[2] *The University of Winnipeg*

[3] *Universite Laval, Quebec, Canada*

Contact: prateekv@alum.iisc.ac.in

Classical approaches to matrix function theory — i.e., extending scalar functions to matrices — are largely organized around two frameworks: entrywise calculus via the Schur (Hadamard) product, and functional calculus via the spectral theorem. In this talk, I present a third, fundamentally different framework based on matrix convolution, in which convolution itself is viewed as a matrix multiplication and gives rise to a new class of matrix transforms.

This perspective leads to a systematic theory of convolution-based transforms and their positivity preserving properties. I will describe convolutional analogues of the classical results of Pólya–Szegő, Schoenberg, Rudin, Loewner, and Horn, including a Schoenberg-type classification theorem showing that absolute monotonicity is again the correct notion for positivity preservers across all dimensions, together with refined results in fixed matrix dimensions.

Beyond positivity, I will highlight two structural features that distinguish convolution from both standard and entrywise products: a Cayley–Hamilton-type theorem for convolution with optimal degree, and a novel polynomial–matrix identity underlying the associated transforms. If time permits, I will also briefly describe an unexpected connection with the Bruhat order on the symmetric group, and an application to sums of discrete random variables, illustrating new links between matrix analysis, algebraic combinatorics, and probability.

(Based on joint work with Javad Mashreghi and Mostafa Nasri.)

On the stability criteria via finite Hankel matrices for regular matrix polynomials

Author: Xuzhou Zhan[1]

[1] *Beijing Normal University at Zhuhai*

Contact: xzzhan@bnu.edu.cn

This talk focuses on several stability criteria via Markov parameters for regular matrix polynomials, which generalize the corresponding criteria constrained by the monic assumption. The testing framework employs two finite Hankel matrices, whose rectangular blocks are the submatrices of the Markov parameters redefined through a column-wise splitting and column reduction for matrix polynomials. Specifically, a Hurwitz stability criterion is characterized by the Hermitian positive definiteness of two finite Hankel matrices. Further, our stability criterion is derived in terms of the Hermitian nonnegative definiteness and nullity of these finite Hankel matrices.

Normal Matrices

Author: Fuzhen Zhang[1]

[1] *Nova Southeastern University*

Contact: zhang@nova.edu

Normal matrices form a central class in matrix analysis, including Hermitian, skew-Hermitian, and unitary, positive semidefinite, permutation matrices and so on. This presentation surveys fundamental properties of normal matrices, including spectral characterization, unitary diagonalization, and trace (in)equality through majorization. It highlights equivalent conditions for normality, with discussions extending to matrix exponentials and polynomials. Examples and counterexamples are provided to clarify certain subtle points about matrix normality.

MS 07: Spectral Graph Theory

Minisymposium Organizers: Sebastian Cioaba, Gabriel Coutinho, Michael Tait

At the request of the organizers, we propose to organize a minisymposium on spectral graph theory. This session will focus broadly on recent developments in spectral graph theory. We expect to invite a speaker list that includes both junior and senior people and that is globally diverse. Spectral graph theory has seen a surge of activity in the last decade, driven both by classical extremal problems and by new techniques from algebra, geometry, and the theory of polynomials. On the spectral extremal side, recent progress has sharpened our understanding of how eigenvalues constrain structure: for instance, new bounds on spectral radius, signless Laplacian extremal problems, and eigenvalue interlacing phenomena have led to refined characterizations of near-extremal graphs. Another major theme driving current research is the inverse eigenvalue problem for graphs, which seeks to understand which spectra arise from graphs with a prescribed structure, or conversely, which structures force spectral constraints. These questions have deep connections to matrix analysis, combinatorics, and theoretical computer science. Finally, spectral graph theory has become increasingly intertwined with discrete geometry, operator theory, and the study of graph-associated polynomials. Techniques involving interlacing polynomials, stability theory, and connections to Euclidean and spherical embeddings have provided powerful unifying frameworks. We will invite speakers in these three broad areas of spectral graph theory.

List of Talks

$q(G)$ for Threshold Graphs

Authors: Luiz Emilio Allem[1], Lucas Siviero Sibemberg[2], Luciano Grippo[3]

[1] *Universidade Federal do Rio Grande do Sul*

[2] *UFRGS*

[3] *Instituto de Ciencias - UNGS - CONICET*

Contact: emilio.allem@ufrgs.br

We study the minimum number of distinct eigenvalues $q(G)$ of threshold graphs. It is known that every threshold graph satisfies $q(G) \leq 4$, which suggests that a complete combinatorial characterization according to the values $q(G) = 2, 3, 4$ is possible. In this talk, we present ongoing results toward such a characterization.

The factorization norm and an inverse theorem for MaxCut

Authors: Igor Balla[1], Istvan Tomon[2], Lianna Hambardzumyan[3]

[1] *Leipzig University*

[2] *Umeå university*

[3] *University of Copenhagen*

Contact: iballa1990@gmail.com

In this talk, we will present a proof of the fact that Boolean matrices with bounded γ_2-norm or bounded normalized trace norm must contain a linear-sized all-ones or all-zeros submatrix, verifying a conjecture of Hambardzumyan, Hatami, and Hatami. We will also discuss further structural results about Boolean matrices of bounded γ_2-norm and applications in communication complexity, operator theory, spectral graph theory, and extremal combinatorics.

As a key application, we establish a theorem for MaxCut which contrasts a celebrated result of Edwards. In particular, we show that if the MaxCut of a graph with m edges is at most $m/2 + O(\sqrt{m})$, then it must contain a clique of size $\Omega(\sqrt{m})$.

Spectral theory of K_t-decomposable graphs

Authors: Matt Burnham[1], Steve Butler[1]

[1] *Iowa State University*

Contact: mattmath@iastate.edu

Given a real number q, the **q-Laplacian** of a graph G is the matrix $A+qD$ where A is the adjacency matrix and D the diagonal degree matrix. If the edge set of G can be partitioned into edge-disjoint copies of K_t, then G is called **K_t-decomposable**.

In this talk, we generalize some results from a survey paper of Cvetković, Rowlinson, and Simić about the signless Laplacian $A+D$ to the $\frac{1}{t-1}$-Laplacian of K_t-decomposable graphs. In particular, we generalize the correspondence between the signless Laplacian and line graphs, positive semi-definiteness, and characterization of graphs with a zero eigenvalue.

Connectivity of distance-regular graphs

Authors: Isabel Byrne[1], Sebastian Cioaba, Jack Koolen[2]

[1] *University of Delaware*

[2] *University of Science and Technology, China*

Contact: itb@udel.edu

In 1985, Brouwer and Mesner proved that the vertex-connectivity of a strongly regular graph equals its valency and the only disconnecting sets of this size are point neighborhoods. In 2009, Brouwer and Koolen generalized this result to distance-regular graphs. In 1996, Brouwer conjectured that the minimum size of a disconnecting set of vertices whose removal disconnects a connected strongly regular graph into non-singleton components equals the size of the neighborhood of an edge. In 2014, Cioaba, Kim, and Koolen disproved Brouwer's conjecture, but also showed that the conjecture is true for many families of strongly regular graphs. In their 2016 survey, van Dam, Koolen, and Tanaka asked whether Brouwer's conjecture is true for distance-regular graph with diameter at least three. In this talk, I will describe our progress regarding this restricted vertex-connectivity of distance-regular graphs.

Cycles in directed graphs

Authors: John Byrne[1], Michael Tait[2]

[1] *University of Delaware*

[2] *Villanova University*

Contact: byrnej37@gmail.com

Let F be a digraph. What is the largest possible minimum outdegree on an n-vertex digraph which does not contain a copy of F? I will discuss algebraic approaches to this question and present some of my related work (joint with Michael Tait).

The non-existence of Moore polygons and spectral Moore bounds

Authors: Hiroshi Nozaki, Sebastian Cioaba, Vishal Gupta[1], Ziqing Xiang

[1] *University of Rochester*

Contact: vishalgupta@rochester.edu

The second largest eigenvalue of a graph's adjacency matrix captures important structure and spectral properties including connectivity and expansion. From the Alon-Boppana theorem, we know that if $\theta < 2\sqrt{k-1}$, then there are only finitely many k-regular graphs with the second largest eigenvalue at most θ. This motivates the following natural question posed by Richey, Stover, and Shutty in 2013. Given $k \geq 3$ and $\theta < 2\sqrt{k-1}$, what is the maximum number of vertices $v(k,\theta)$ of a k-regular graph G with $\lambda_2(G) \leq \theta$? This problem can be seen as a spectral version of the Moore (or degree-diameter) problem. In this talk, I will present recent progress on this question. This is joint work with Cioabă, Nozaki, and Xiang.

The inverse symplectic eigenvalue problem of a graph

Author: Leslie Hogben[1]

[1] *American Institute of Mathematics, Iowa State University, Purdue University*

Contact: hogben@aimath.org

The graph $\mathcal{G}(A)$ of a real symmetric matrix $n \times n$ matrix $A = [a_{ij}]$ has vertices $V = \{1, \ldots, n\}$ and edges $E = \{\{i,j\} : a_{ij} \neq 0 \text{ and } i \neq j\}$. The set of matrices described by a graph G is $\mathcal{S}(G) = \{A \in \mathbb{R}^{n\times n} : \mathcal{G}(A) = G \text{ and } A^\top = A\}$ and the Inverse Eigenvalue Problem of G (IEP-G) is to determine all possible spectra (i.e., multisets of eigenvalues) of matrices $A \in \mathcal{S}(G)$. Other inverse eigenvalue problems for G have also been studied based on matrix properties, e.g., the PSD Inverse Eigenvalue Problem of G where matrices must be positive semidefinite.

A symplectic matrix is a $(2n) \times (2n)$ real matrix S such that $S^\top \Omega S = \Omega$ where $\Omega = \left[\begin{array}{c|c} O & I \\ \hline -I & O \end{array}\right]$. It is known that for any $(2n) \times (2n)$ symmetric positive definite matrix A, there exists a symplectic matrix S and an $n \times n$ diagonal matrix D such that $S^\top A S = \left[\begin{array}{c|c} D & O \\ \hline O & D \end{array}\right]$. The diagonal entries of D are unique (up to re-ordering) and are called the symplectic eigenvalues of A; the symplectic spectrum of A is the multiset of symplectic eigenvalues of A. The Inverse Symplectic Eigenvalue Problem of G (ISEP-G) is to determine all possible symplectic spectra of a labeled graph G with vertex set $\{1, \ldots, 2n\}$ (unlike the IEP-G, the labeling affects the symplectic spectrum). There are many differences between the IEP-G and the ISEP-G, e.g., that there are many labelled graphs that allow all possible symplectic spectra including all symplectic eigenvalues equal. Whereas, for the IEP-G, all eigenvalues equal is allowed only by a graph with no edges.

This talk will provide an introduction to the ISEP-G, including methods to construct matrices with any given symplectic spectrum and solutions to the ISEP-G for all labeled graphs of order four.

A graph energy conjecture through the lenses of semidefinite programming

Authors: Emanuel Juliano[1], Aida Abiad[2], Gabriel Coutinho

[1] *Federal University of Minas Gerais*

[2] *Eindhoven University of Techonolgy*

Contact: emanuelsilva@dcc.ufmg.br

In the 1980s, using the software Graffiti, Fajtlowicz made a conjecture relating the independence number of a graph with its energy, a spectral parameter introduced by Gutman (1978). By formulating the graph energy as a semidefinite program (SDP), we take a step towards Fajtlowicz's conjecture, relating the graph energy to the fractional clique covering number. As a byproduct of the SDP formulation, we obtain several lower bounds for the graph energy that improve and refine previous results by Hoffman (1970) and Nikiforov (2007).

Delsarte designs and Galois groups

Authors: Jesse Lansdown[1], William Martin[2]

[1] *University of Galway*

[2] *Worcester Polytechnic Institute*

Contact: martin@wpi.edu

Delsarte theory has been extended to the study of subsets in an increasing variety of association schemes in recent years, with many different motivations and applications. Tools originally developed for the study of error-correcting codes in the Hamming scheme and combinatorial t-designs in the Johnson scheme apply equally well in association schemes with irrational eigenvalues. The goal of the talk is to study the Delsarte T-designs in an assocaition scheme with irrational eigenvalues. We show that, for any subset T of eigenspaces, a T-design must also be a T'-design where T' the closure of T under the action of the Galois group of the splitting field of the association scheme, acting on primitive idempotents. Conjugacy class schemes of finite groups will be the main examples explored. This is based on joint work with Jesse Lansdown (Galway).

How many eigenvalues of a tree can attain the maximum multiplicity?

Authors: Tracy Hall[1], Rupert Levene[2], Seth Meyer, Shahla Nasserasr[3], Polona Oblak[4], Helena Šmigoc[2]

[1] *Hall Labs LLC*

[2] *University College Dublin*

[3] *Rochester Institute of Technology*

[4] *University of Ljubljana*

Contact: polona.oblak@fri.uni-lj.si

For a tree T we consider the set $\mathcal{S}(T)$ of real symmetric matrices whose off-diagonal zero-nonzero pattern is equal to the pattern of the adjacency matrix of T. It is well known that the maximum multiplicity of an eigenvalue over matrices in $\mathcal{S}(T)$ is equal to the path cover number $P(T)$ of the tree T.

We present a novel decomposition of the tree into a set of vertices and paths, which serves as a tool for analysing matrices achieving maximum multiplicity $P(T)$. We define a parameter $\mathrm{MM}(T)$ to be the maximal positive integer k for which there exists a matrix $A \in S(T)$ with k eigenvalues achieving the multiplicity $P(T)$. We derive a combinatorial upper bound for $\mathrm{MM}(T)$ and identify some families of trees that achieve this bound.

Zero-Nonzero Patterns of Symmetric Orthogonal Matrices

Authors: Wayne Barrett[1], Shaun Fallat[2], Veronika Furst[3], Shahla Nasserasr[4], Brendan Rooney[4], Michael Tait[5]

[1] *BYU*

[2] *University of Regina*

[3] *Fort Lewis College*

[4] *Rochester Institute of Technology*

[5] *Villanova University*

Contact: brsma@rit.edu

Given a graph G on n vertices, $\mathcal{S}(G)$ is the set of symmetric $n \times n$ matrices with the same off-diagonal zero pattern as the adjacency matrix $A(G)$. We say that a graph G has $q(G) = 2$ if there is a matrix $M \in \mathcal{S}(G)$ with exactly 2 distinct eigenvalues. This is equivalent to the existence of an orthogonal matrix $M \in \mathcal{S}(G)$. We are interested in understanding the graphs G for which $q(G) = 2$ (or equivalently, the zero-nonzero patterns of symmetric orthogonal matrices).

In this talk, we focus on the following conjecture: if the complement of G has at most $n - 3$ edges, then $q(G) = 2$. This conjecture has been resolved for graphs with bipartite complement,

and verified in some cases where the complement of G is not bipartite. We report on this ongoing project.

Semidefinite programming bounds on fractional cut-cover and maximum 2-SAT for highly regular graphs

Authors: Gabriel Coutinho, Henrique Soares Assumpção e Silva[1]

[1] *Universidade Federal de Minas Gerais*

Contact: henriquesas2010@gmail.com

We use semidefinite programming to bound the fractional cut-cover parameter of graphs in association schemes in terms of their smallest eigenvalue. We also extend the equality cases of a primal-dual inequality involving the Goemans-Williamson semidefinite program, which approximates MAXCUT, to graphs in certain coherent configurations. Moreover, we obtain spectral bounds for MAX 2-SAT when the underlying graphs belong to a symmetric association scheme by means of a certain semidefinite program used to approximate quadratic programs, and we further develop this technique in order to explicitly compute the optimum value of its gauge dual in the case of distance-regular graphs.

MS 08: The Inverse Eigenvalue Problem of a Graph and Zero Forcing

Minisymposium Organizers: Leslie Hogben, Bryan Shader

Inverse problems play a central role in mathematics and arise naturally in applications. In many instances the inverse problem reduces to a question about the existence of a matrix with a prescribed set of eigenvalues and prescribed structure. The Inverse Eigenvalue Problem of a Graph (IEP-G) studies such questions, where the prescribed structure is the matrix is symmetric and there is a fixed nonzero off-diagonal pattern; this pattern is described by the edges of a graph. Due to the lack of effective tools for the IEP-G at the time, much early research focused on subquestions that could lead to progress on the IEP-G. The most important of these is the study of the maximum possible multiplicity of an eigenvalue among matrices described by G, or equivalently, the maximum possible nullity. The zero forcing number arose independently in several applications, including as an upper bound for the maximum nullity of a graph. For both these problems, numerous variants are now studied and there have been important recent developments, especially the introduction of additional strong properties for the IEPG, the introduction of propagation time and throttling for zero forcing, and the study of nodal domains of discrete Schrodinger operators. Talks in this mini-symposium will report on recent developments.

List of Talks

A unified framework for the Expander Mixing Lemma for graphs and its applications

Authors: Aida Abiad[1], Sjanne Zeijlemaker[1]

[1] *Eindhoven University of Techonolgy*

Contact: a.abiad.monge@tue.nl

A unified framework of the Expander Mixing Lemma for irregular graphs using adjacency eigenvalues will be presented, as well as several new versions of it. We will also show some of its applications in graph theory, which include spectral bounds on the zero forcing number of a graph. To derive our results we use a new application of weight partitions of graphs, where the Perron eigenvector entries are used as vertex weights as a way to regularise the graph.

An inverse eigenvalue problem for structured matrices determined by graph pairs

Authors: Adam Berliner[1], Michael Cavers[2], Minerva Catral[3], Pauline van den Driessche[4], Sooyeong Kim[5]

[1] *St Olaf College*

[2] *University of Toronto*

[3] *Xavier University*

[4] *University of Victoria*

[5] *University of Guelph*

Contact: catralm@xavier.edu

For vertex-labelled graphs G and H on $n \geq 1$ vertices, we consider matrices of the form $C(A,B) = \left[\begin{array}{c|c} A & B \\ \hline I & O \end{array}\right] \in \mathbb{R}^{2n\times 2n}$ where $A, B \in \mathbb{R}^{n\times n}$ are a pair of real symmetric matrices with nonzero patterns determined by the edges of the graph pair G, H. We denote the set of all such matrices by $\mathcal{S}(G,H)$. Our aim is to determine all possible spectra for $C(A,B) \in \mathcal{S}(G,H)$. We conjecture that $C(A,B)$ can attain any spectrum invariant under conjugation regardless of the chosen vertex-labelled graphs G and H. In this talk, we highlight some results that support our conjecture.

Resolving Inverse Singular Value Problems with Spoiler Spaces

Authors: Bryan Shader[1], Caleb Cheung[1]

[1] *University of Wyoming*

Contact: ccheung@uwyo.edu

Let A be an $m \times n$ matrix. The spoiler space of A is the set of all $m \times n$ matrices X such that, $XA^\top$ is symmetric, $A^\top X$ is symmetric, and $X \circ A = O$, where "$\circ$" denotes the entrywise Schur product. If the spoiler space of A contains only the 0 matrix, we say that A has the Strong Singular Value Property (SSVP). The SSVP gives us access to a rich toolset for resolving inverse singular value problems. In this talk, we will discuss results pertaining to the spoiler space, its orthogonal complement, the tangent space, as well as the singular value multiplicity lists of various matrices and how they are connected.

Inverse Eigenvalue Problems for Graphs: The Weighted Laplacian Case

Authors: Shaun Fallat[1], Himanshu Gupta[1], Jephian C.-H. Lin[2], Minerva Catral[3]

[1] *University of Regina*

[2] *National Yang Ming Chiao Tung University*

[3] *Xavier University*

Contact: shaun.fallat@uregina.ca

An inverse eigenvalue problem for a graph (IEP-G) asks a fundamental question: What are the possible spectra for (symmetric) real matrices fitting a given graph? Many have worked on several aspects of the IEP-G with exciting advances and variations appearing over the past forty years. Here, the focus will be on weighted Laplacian matrices associated with a graph. Such matrices are permanently intertwined in combinatorial matrix theory and numerical analysis and form a natural matrix class to study for an inverse eigenvalue problem. For a given graph G, we aim to determine the possible realizable spectra for a generalized (or weighted) Laplacian matrix associated with G. We present such results for certain families of graphs and graphs on a small number of vertices, including possible ordered multiplicity lists. Strong matrix properties have been associated with several adaptations of inverse eigenvalue problems and sign pattern eigenvalue problems. We introduce a new strong property, the strong spectral property for weighted Laplacian matrices (SSPWL), and establish its Supergraph and Bifurcation lemmas. We develop a Jacobian method which can be used to verify such a strong property. These tools are then applied to derive potential spectral regions of weighted Laplacian matrices of a graph on four vertices.

Parameters connected to the strong nullity interlacing property

Authors: Mary Flagg[1], Aida Abiad[2], Bryan Shader[3], Jephian Lin[4], Tracy Hall[5]

[1] *University of St. Thomas*

[2] *Eindhoven Institute of Technology*

[3] *University of Wyoming*

[4] *National Sun Yat-sen University*

[5] *Hall Labs LLC*

Contact: flaggm@stthom.edu

Given a simple graph G, $\mathcal{S}(G)$ is the set of real symmetric matrices indexed by the vertices in G and with off-diagonal zeros corresponding to non-edges in G. The problem of finding the maximum nullity of a matrix in $\mathcal{S}(G)$ has been extensively studied. We consider the maximum nullity of a matrix A and its principal submatrix $A(i)$ corresponding to deleting the vertex i. The strong Arnold property possessed by some matrices is one of the useful tools for studying maximum nullity. The parameter $\xi(G)$ is the maximum nullity of a matrix in $\mathcal{S}(G)$ with the strong Arnold property, and has been shown to be minor monotone. We defined a new parameter which we call $\xi\xi(G)$ for the maximum sum of the nullities of A and $A(i)$ for matrices $A \in \mathcal{S}(G)$. We show that $\xi\xi(G)$ is also minor monotone. In this talk I will show how $\xi\xi(G)$ is a refinement of $\xi(G)$ and minimal minors for small values of $\xi\xi(G)$.

The Inverse Symplectic Eigenvalue Problem and Coupled Zero Forcing for Graphs

Authors: Bryan Shader[1], Himanshu Gupta[2], Leslie Hogben[3], Tony Wong[4]

[1] *University of Wyoming*

[2] *University of Regina*

[3] *American Institute of Mathematics, Iowa State University, Purdue University*

[4] *Kutztown University of Pennsylvania*

Contact: himanshu.gupta@uregina.ca

Symplectic geometry appears in many areas of mathematics, physics, and applications, and naturally gives rise to interesting matrix families and properties. Symplectic eigenvalues extend the classical notion of eigenvalues to the symplectic setting and are guaranteed to exist for positive definite matrices by Williamson's theorem. We introduce the inverse symplectic eigenvalue problem for positive definite matrices whose zero and nonzero pattern is described by a labeled graph (ISEPG). In this talk, we define the ISEPG and present key tools developed to address it. We focus particularly on coupled graph zero forcing, a combinatorial technique used to bound the maximum symplectic eigenvalue multiplicity for a given graph. This is joint work with Leslie Hogben, Bryan Shader, and Tony Wong.

A general strong property for IEP-G

Author: H. Tracy Hall
Contact: h.tracy@gmail.com

The Inverse Eigenvalue Problem for a Graph (IEP-G) asks for the possible spectra of a real symmetric matrix knowing only which off-diagonal entries are non-zero, as described by a graph G. Three matrix properties, collectively called the "strong properties", have become prominent in the study of this problem, due in part to their good behavior with respect to edge deletion and contraction within G. The three properties differ in what eigenvalue data is preserved:

- For the Strong Arnold Property (SAP), a single multiplicity.
- For the Strong Multiplicity Property (SMP), all multiplicities.
- For the Strong Spectral Property (SSP), all multiplicities and locations.

In most cases, a question that calls for a strong property will be served well by one of these three perspectives. There does, however, exist a more general perspective that contains each of these as a special case, here introduced as the General Strong Property (GSP). For the GSP, any chosen subset of eigenvalue multiplicities can be preserved, and eigenvalue locations can be allowed to vary locally subject to any set of linear constraints. The usual strong property consequences follow, which include: an algebraic definition, a verification matrix, superpatterns, edge decontraction, matrix liberation, and bifurcation.

In addition, the null space characterization of SAP given by H. van der Holst is extended to an eigenspace characterization of GSP, including the special cases of SMP and SSP, whose eigenspace characterization has not appeared in the literature.

Tools for Determining the Minimum Rank of a Graph

Author: Mark Hunnell[1]

[1] *Winston-Salem State University*

Contact: hunnellm@wssu.edu

The minimum rank of a graph G of order n is the smallest possible rank over all real symmetric $n \times n$ matrices A whose (i, j)th entry, for $i \neq j$, is nonzero whenever ij is an edge of G and zero otherwise. We discuss some refinements of techniques currently in the literature to determine the minimum rank of a graph, some new tools to bound this value, and an approach for understanding the gaps in values between parameters used to bound the minimum rank of a graph. We also discuss implementations of known techniques, algorithmic improvements, and applications in computer assisted experimentation for the minimum rank problem.

Orientable forcing and relationships with linear algebra

Authors: Bonnie Jacob[1], Erika L.C. King[2]

[1] *Rochester Institute of Technology*

[2] *Hobart and William Smith Colleges*

Contact: bcjntm@rit.edu

In this talk, we introduce a new parameter, the *orientable forcing number* of an undirected graph G, which is the maximum zero forcing number among all oriented graphs that have G as their underlying undirected graph. We establish some properties of the orientable forcing number, including extreme values, and discuss how the parameter relates to matrices.

Graph Products to Achieve few Distinct Eigenvalues

Author: Mark Kempton[1]

[1] *Brigham Young University*

Contact: mkempton@mathematics.byu.edu

One of the most challenging aspects of graph inverse eigenvalue problems is knowing when a graph admits a matrix with few distinct eigenvalues. We will discuss results where we construct matrices with only two distinct eigenvalues corresponding to graphs that arise from various kinds of product structures. We will discuss also some directions for future work along these lines.

Jacobian method and strong properties

Authors: Minerva Catral[1], Shaun Fallat[2], Himanshu Gupta[2], Jephian C.-H. Lin[3]

[1] *Xavier University*

[2] *University of Regina*

[3] *National Yang Ming Chiao Tung University*

Contact: jephianlin@gmail.com

A sign pattern is a matrix whose entries are in $\{+, -, 0\}$, while its quantitative class is the set of real matrices whose entries match the corresponding signs. A sign pattern is said to be spectrally arbitrary if its quantitative class contains matrices demonstrating all possible monic real polynomials as the characteristic polynomials. Historically, there are the Jacobian method, the nilpotent-centralizer method, and the nonsymmetric strong spectral property (SSP, which is equivalent to the similarity-transversality property) to witness a spectrally arbitrary sign pattern. However, their

relations are unclear. In this talk, we survey these methods and show that they are equivalent when the minimal polynomial has degree n.

Leaky Forcing of Unicyclic Graphs

Authors: Carolyn Reinhart[1], Beth Bjorkman[2], Franklin Kenter[3], Lei Cao[4], Ryan Moruzzi[5], Violeta Vasilevska[6]

[1] *Swarthmore College*

[2] *Air Force Research Laboratory Sensors Directorate*

[3] *United States Naval Academy*

[4] *Nova Southeastern University*

[5] *California State University, Northridge*

[6] *Utah Valley University*

Contact: creinha1@swarthmore.edu

Zero forcing is a graph coloring process in which a set of initially blue vertices force the remaining vertices in the graph to be colored blue after repeated applications of a color change rule. Leaky forcing is a fault-tolerant variant of zero forcing in which some set of ℓ vertices, called leaks, are forbidden from forcing. The ℓ-leaky forcing number is the size of the smallest set of initially blue vertices for which the graph is successfully forced despite the placement of ℓ leaks. In this talk, we will characterize the ℓ-leaky forcing number for unicyclic graphs.

A Graphical Approach to Isospectral Unfoldings

Authors: Dallin Seyfried[1], Mark Kempton[1]

[1] *Brigham Young University*

Contact: dseyfr99@gmail.com

An isospectral reduction is a method of shrinking a large matrix into a smaller matrix while preserving properties of the original's spectrum. The inverse, an isospectral unfolding, takes a matrix of an isospectral reduction and expands it into a larger matrix that has that reduction. We present a system of nonlinear equations forming the foundation of general isospectral unfolding. Graphs naturally tie in to isospectral reductions via their adjacency matrices. Thus, automorphic orbits, equitable partitions, and their relations to graphs sharing the same isospectral reductions will be explored, theorems given, and proofs discussed. Connections to quantum walk matrices will also be made.

Sign patterns that require or allow the non-symmetric strong spectral property

Author: Kevin Vander Meulen[1]

[1] *Redeemer University*

Contact: kvanderm@redeemer.ca

A matrix A has the non-symmetric strong spectral property (nSSP) if $X = O$ is the only matrix which satisfies $A \circ X = O$ and $AX^T = X^T A$. This property comes with implications for eigenvalue properties of sign patterns, including a bifurcation lemma and superpattern lemma. We describe some classes of sign patterns for which every matrix with the sign pattern will have the nSSP. We also provide methods for recognizing when a sign pattern restricts matrices from having the nSSP. In particular, we provide a characterization of the sign patterns that allow a matrix with the nSSP. This presentation is based on joint work with Michael Cavers and Zhongshan Li.

MS 09: Numerical Ranges and Numerical Radii

Minisymposium Organizers: Pan Shun Lau, Chi-Kwong Li, Raymond Nung-Sing Sze

The study of numerical ranges and numerical radii has a long and distinguished history, tracing back to the early 20th century. Over the decades, it has yielded substantial and influential results, with applications across both pure and applied mathematics, including operator theory, functional analysis, matrix inequalities, perturbation theory, numerical analysis, and quantum information science. The goal of this mini-symposium is to stimulate research and foster interaction among scholars working on numerical ranges and radii, their generalizations, and related areas.

List of Talks

On the k-numerical ranges of matrices

Author: Kennett Dela Rosa[1]

[1] *University of the Philippines Diliman*

Contact: prkdelarosa@math.upd.edu.ph

This study considers some problems involving the k-numerical range. Following the idea of the zero-dilation index, the notion of the zero-trace index is introduced, which is defined as the largest zero-trace compression of a matrix. Alternative characterization of the zero-trace index is given, and zero-trace indices of certain classes of matrices are identified. The study also considers recent results on the numerical range of cyclic shift matrices. A recent solution to a conjecture proved that a certain arrangement of the weights of a given cyclic shift matrix maximizes the classical numerical

range. This work explores whether analogous optimal arrangements exist for the k-numerical range.

Norms on tensors in quantum information related to numerical radii

Author: Shmuel Friedland[1]

[1] *University of Illinois Chicago*

Contact: friedlan@uic.edu

In this talk we survey some recent results on entanglement and separability of general, symmetric (bosons), skew-symmetric (fermions) tensors, and their computability. All these results are related to corresponding numerical radii. This talk is based on the arXiv preprint

S. Friedland, Tensors, entanglement, separability, and their complexity, arXiv:2509.21639, 2025.

The Cauchy Equation and Norm-Additive Mappings on Positive Cones of Commutative C^*-Algebras

Author: Daisuke Hirota[1]

[1] *National Institute of Technology, Tsuruoka College*

Contact: dhirota@tsuruoka-nct.ac.jp

The Cauchy functional equation plays a fundamental role in the study of additive and linear structures arising from numerical and norm-related information in functional analysis. In this talk, we investigate preserver problems on positive cones of commutative C*-algebras, where a norm identity of Fischer–Muszély type, arising from the Cauchy functional equation, determines the underlying algebraic and geometric structure.

Let A_i $(i = 1, 2)$ be commutative C*-algebras, and let A_i^+ denote their positive cones. We consider surjective mappings $T : A_1^+ \to A_2^+$, not assumed to be continuous, satisfying the norm identity

$$\|T(a + b)\| = \|T(a) + T(b)\| \qquad (a, b \in A_1^+).$$

We show that every such mapping is necessarily additive and positive homogeneous. Moreover, in the unital case, if T is injective, then it can be normalized to a composition operator, which in turn induces an isometric isomorphism between the underlying commutative C*-algebras A_1 and A_2.

These results demonstrate that norm identities of Fischer–Muszély type on positive cones contain rich numerical and order-theoretic information, which completely determines the structure of

commutative C*-algebras.

On the ellipticity of the higher rank numerical range

Authors: Rute Lemos[1], Natália Bebiano[2], Graça Soares[3]

[1] *CIDMA, University of Aveiro, Portugal*

[2] *CMUC, University of Coimbra*

[3] *CMAT-UTAD, University of Trás-os-Montes e Alto Douro*

Contact: rute@ua.pt

The higher rank numerical range is investigated for 2-by-2 block matrices with associated Kippenhahn curves consisting of ellipses and eventually points. As a consequence, elliptical higher rank numerical range results are derived in a unified way, using an approach developed by Spitkovsky et al.

Linear maps preserving products of involutions

Authors: Tejbir Lohan[1], Chi-Kwong Li[2], Sushil Singla[3]

[1] *Indian Statistical Institute, Delhi Centre, India*

[2] *College of William & Mary, USA*

[3] *University of Regina, Canada*

Contact: tejbirlohan70@gmail.com

Linear preserver problems study linear maps on matrix spaces that leave certain functions, subsets, or relations invariant, whereas matrix decomposition problems focus on expressing matrices as products of matrices with prescribed structural properties. An element of the algebra $M_n(\mathbb{F})$ of $n \times n$ matrices over a field $\mathbb{F}$ is called an involution if its square equals the identity matrix. A classical result of Gustafson, Halmos, and Radjavi asserts that any product of involutions in $M_n(\mathbb{F})$ can be written as a product of at most four involutions. It is also known that a matrix is a product of two involutions if and only if it is similar to its inverse. These results naturally lead to the following linear preserver question: which linear maps on $M_n(\mathbb{F})$ preserve products of involutions?

In this talk, we address this question and present a characterization of bijective linear maps that preserve matrices expressible as products of two, three, or four involutions in $M_n(\mathbb{F})$. We also outline possible directions for further research. This is joint work with Chi-Kwong Li and Sushil Singla.

Multiplicities and k-Numerical Range

Authors: Nancy Menzelthe[1], Pan Shun Lau[1], Tin-Yau Tam[1]

[1] *University of Nevada, Reno*

Contact: nmenzelthe@unr.edu

Given $1 \leq k \leq n$, the k-numerical range of $A \in \mathbb{C}_{n\times n}$ is defined by

$$W_k(A) := \left\{ \sum_{i=1}^{k} x_i^* A x_i : x_1, \ldots, x_k \text{ orthonormal vectors in } \mathbb{C}^n \right\} \subset \mathbb{C}.$$

Motivated by Davis' intuitive explanation of the Elliptical Range Theorem, we introduce two notions of multiplicity for points in $W_k(A)$, namely wedge multiplicity and projection multiplicity. The wedge multiplicity is related to the Grassmannian and the projection multiplicity is related to the set of rank k orthogonal projectors. We present several results concerning each notion and provide examples illustrating these multiplicities. The corresponding real analogues $V_k(A)$ are also studied.

Matrix numerical ranges of Toeplitz operators with polynomial symbols

Author: Linda Patton[1]

[1] *Cal Poly San Luis Obispo*

Contact: lpatton@calpoly.edu

Using results from Brown-Halmos, Klein (1972) described the numerical range of a general Toeplitz operator on $H^2(\mathbb{D})$. In particular, the numerical range of a Toeplitz operator T_p with polynomial symbol p is the convex hull of the image of the unit disk under p. By analyzing $p(\mathbb{T})$ and its relationship to the Kippenhahn curve of a matrix, we provide conditions under which the closure of $W(T_p)$ is the numerical range of a finite matrix M_p. Results about the eigenvalues of M_p will be discussed as well.

On the spatial numerical range

Author: Edward Poon
Contact: poon3de@erau.edu

The spatial numerical range of an operator T on a normed space $(\mathcal{X}, \|\cdot\|)$ is the set

$$W(T) = \{f(Tx) : x \in \mathcal{X}, f \in \mathcal{X}^*, \|x\| = \|f\|^d = f(x) = 1\};$$

when $\|\cdot\|$ is induced by an inner product this coincides with the classical numerical range. We investigate some properties of the spatial numerical range.

Numerical Range of Positive Hermitian Hankel Matrices

Author: Brooke Randell[1]

[1] *UCSC*

Contact: brandell@ucsc.edu

In 2022, Hwa-Long Gau and Pei Yuan Wu discussed the numerical range of various Hankel matrices with an emphasis on which subsets of the complex plane are attainable. We will be discussing as well as expanding upon these results when the Hankel matrices are positive and Hermitian.

MS 10: Inverse Problems and Uncertainty Quantification through the Lens of Numerical Linear Algebra

Minisymposium Organizers: Jonathan Lindbloom, Toluwani Okunola, Mirjeta Pasha

Inverse problems arise in all areas of science, engineering, and industry, enabling the recovery of hidden parameters or structures from indirect and often noisy data. Their solution—and the associated uncertainty quantification—rely on computational algorithms in which numerical linear algebra plays a central role. This mini-symposium brings together researchers working on a variety of theoretical and applied aspects of inverse problems and UQ, with a particular emphasis on the underlying NLA tools (including software) that expose problem structure and enable scalable, reliable inference.

List of Talks

Time-varying Bayesian Inverse Problems with Sparse Priors and Randomization

Authors: Diego Arenas Mata[1], Mirjeta Pasha

[1] *Virginia Tech*

Contact: camiloam@vt.edu

Sparse priors, such as the Laplace prior, are of considerable interest in Bayesian inverse problems because they promote sparsity and preserve edges in the solution, which are often more appropriate than the smooth reconstructions obtained with Gaussian priors. However, sampling from the resulting non-Gaussian posteriors is challenging, particularly in high-dimensional settings. To address this, we build on the randomize-then-optimize (RTO) framework and its extension to ℓ_1-type priors via variable transformation. We propose a method for posterior sampling in Bayesian inverse problems with Laplace priors that converts the prior into a standard Gaussian in the transformed space. Within the RTO framework, this results in a nonlinear optimization problem at each iteration, which we solve using the Levenberg–Marquardt algorithm. To accelerate the linear system solves at each Levenberg–Marquardt iteration, we employ randomized preconditioners and investigate the effect of using single and half precision arithmetic in these solves. Performance is evaluated through numerical experiments on image deconvolution and computed tomography problems.

Interpolation-Based Estimation and Uncertainty Quantification of Periodic Time-Varying Parameters

Authors: Andrea Arnold[1], Sara Johnson[2]

[1] *Worcester Polytechnic Institute*

[2] *University of Puget Sound*

Contact: anarnold@wpi.edu

Many applications in modern day science involve unknown system parameters that must be estimated from limited data. A subset of these problems involves parameters that vary with time but have unknown evolution models and cannot be directly observed. In this work, we formulate time-varying parameter estimation in deterministic dynamical systems as an interpolation problem, where the function values at the node locations are unknown and must be estimated to obtain

a polynomial-based approximation of the unknown parameter. We discuss the numerical considerations that arise when solving the linear systems in these approximations, along with several computed examples.

Stochastic zeroth-order calculation of operator quantities

Authors: Jonas Bresch[1], Garbiele Steidl[1], Martin Schoen[1], Oleh Melnyk[1]

[1] *Technische Universität Berlin*

Contact: bresch@math.tu-berlin.de

The maximization of the (generalized) Rayleigh quotient is a central problem in numerical linear algebra. Conventional algorithms for its computation typically rely on matrix-adjoint products, making them sensitive to errors arising from adjoint mismatches. To address this issue, we introduce a stochastic zeroth-order Riemannian algorithm that maximizes the generalized Rayleigh quotient without requiring adjoint or matrix inverse computations. Moreover, the construction can be translated to other quantities as for example $\|A\|$ and $\|A - V\|$, where only evaluations of the linear map $x \mapsto Ax$, respectively $y \mapsto V^*y$ are available. Those are of interest in inverse problems and uncertainty quantification. We provide theoretical convergence guarantees showing that the iterates converge to the set of global maximizers of the (generalized) Rayleigh quotient at a sublinear rate with probability one. Our theoretical results are supported by numerical experiments, which demonstrate the excellent performance of the proposed method compared to state-of-the-art algorithms.

Streaming Algorithms for Big Data Inverse Problems

Authors: Eric de Sturler[1], Malena Sabate Landman[2], Priyanka Sinha[1]

[1] *Virginia Tech*

[2] *University of Bath*

Contact: sturler@vt.edu

Big data applications are becoming ever more prominent, and in many applications we need to solve very large linear or nonlinear inverse problems while handling only a relatively small amount of data at a time. Moreover, we are interested in distributed, possibly asynchronous, algorithms that solve large problems while only exchanging limited information. We need algorithms that combine approximate (partial) solutions with incoming data or data read from secondary memory to incrementally further improve the solution. We will discuss several algorithmic variations and their convergence.

Projected Regularization in Low Precision

Authors: Chelsea Drum[1], James Nagy[1], Lucas Onisk[1]

[1] *Emory University*

Contact: cndrum@emory.edu

In recent years, mixed-precision and reduced-precision algorithms for solving large-scale linear systems have emerged as an effective approach for exploiting modern GPU architectures. While much of this work has focused on well-conditioned systems, comparatively little attention has been given to ill-posed inverse problems, where regularization is essential. In this talk we consider projected iterated Tikhonov regularization methods in reduced precision and show that these methods can produce reconstructions comparable to their high-precision counterparts. In addition, we discuss a secant-type update for automatic regularization parameter selection within the Golub–Kahan bidiagonalization framework, and demonstrate its effectiveness in the reduced-precision setting.

Separable Nonlinear Bayesian Inverse Problems

Authors: Jordan Dworaczyk[1], Malena Espanol[1]

[1] *Arizona State University*

Contact: malena.espanol@asu.edu

Separable nonlinear inverse problems arise in many applications where a forward model depends linearly on some unknowns and nonlinearly on others, including semi-blind deconvolution. We adopt a Bayesian framework with Gaussian noise and Gaussian priors on the linear variables, leading to regularized formulations of the inverse problem. We examine prior models for the nonlinear parameters and show that maximum a posteriori (MAP) estimation yields regularized separable nonlinear least squares problems that can be efficiently solved using variable projection (VarPro) methods, as demonstrated through numerical examples.

A Scalable Sequential Framework for Dynamic Inverse Problems via Model Parameter Estimation

Authors: Aryeh Keating[1], Mirjeta Pasha

[1] *Virginia Tech*

Contact: aryehkeating@vt.edu

Large-scale dynamic inverse problems are often ill-posed due to model complexity and the high dimensionality of the unknown parameters. Regularization is commonly employed to mitigate

ill-posedness by incorporating prior information and structural constraints. However, classical regularization formulations are frequently infeasible in this setting due to prohibitive memory requirements, necessitating sequential methods that process data and state information online, eliminating the need to form the full space-time problem. In this work, we propose a memory-efficient framework for reconstructing dynamic sequences of undersampled images from computerized tomography data that requires minimal hyperparameter tuning. The approach is based on a prior-informed, dimension-reduced Kalman filter with smoothing. While well suited for dynamic image reconstruction, practical deployment is challenging when the state transition model and covariance parameters must be initialized without prior knowledge and estimated in a single pass. To address these limitations, we integrate regularized motion models with expectation-maximization strategies for the estimation of state transition dynamics and error covariances within the Kalman filtering framework. We demonstrate the effectiveness of the proposed method through numerical experiments on limited-angle and single-shot computerized tomography problems, highlighting improvements in reconstruction accuracy, memory efficiency, and computational cost.

A Provably Convergent MM-GKS Variant for Large-Scale Image Reconstruction Problems

Authors: Misha Kilmer[1], Eric de Sturler[2], Mirjeta Pasha

[1] *Tufts University*

[2] *Virginia Tech*

Contact: misha.kilmer@tufts.edu

Reconstructing high-quality images with sharp edges requires edge-preserving regularization, often imposed using the ℓ_1-norm of the gradient. To get a computationally tractable problem, the ℓ_1-norm term is typically replaced with a sequence of ℓ_2-norm weighted gradient terms with the weights determined from the current solution estimate. The majorization-minimization generalized Krylov subspace method (MM-GKS) has the advantage of combining the updating of the regularization operator with generalized Krylov subspaces (GKS). Unfortunately, the storage and the cost of repeated orthogonalization can present overwhelming memory requirements and computational costs.

We present a variant of MM-GKS that provably converges to the minimum of the smoothed functional even if the solution search space dimension remains very small. This substantially improves theoretical results for MM-GKS where the convergence proof relies on (eventually) spanning the full problem space. Using this result, we develop a new method that solves the minimization/imaging problem by alternatingly compressing and expanding the search space while maintaining strict monotonic convergence. Our method can solve large-scale problems efficiently both in terms of computational complexity and memory requirements. We further generalize our proposed method to handle streaming problems where the data is either not all available simultaneously or the size of the problem demands it be treated as such. We demonstrate the utility of our approach on several

image reconstruction and restoration problems. This is joint work with Mirjeta Pasha (VT) and Eric de Sturler (VT).

Scalable iterative data-adaptive RKHS regularization for linear inverse problems

Authors: Haibo Li[1], Jinchao Feng[2], Fei Lu[3]

[1] *Huazhong University of Science and Technology*

[2] *Great Bay University*

[3] *Johns Hopkins University*

Contact: haibolee1729@gmail.com

In this talk, I will present iDARR, a scalable iterative Data-Adaptive RKHS Regularization method for solving ill-posed linear inverse problems. This method searches for solutions in subspaces where the true solution can be identified, with the data-adaptive reproducing kernel Hilbert space (RKHS) penalizing the spaces of small singular values. At the core of the method is a new generalized Golub-Kahan bidiagonalization procedure that recursively constructs orthonormal bases for a sequence of RKHS-restricted Krylov subspaces. The method is scalable, with a complexity of O(kmn) for m-by-n matrices, where k denotes the number of iterations. Numerical tests on the Fredholm integral equation and 2D image deblurring demonstrate that it outperforms the widely used L^2 and l^2 norms, consistently producing stable and accurate solutions that converge when the noise level decreases.

Multigrid-Accelerated Sparsity-Promoting Projection Methods for Inverse Problems

Authors: Jonathan Lindbloom[1], Mirjeta Pasha, Jan Glaubitz[2]

[1] *Dartmouth College*

[2] *Linköping University*

Contact: jonathan@lindbloom.com

Hybrid projection methods are an effective iterative approach for the solution of large-scale linear inverse problems, including those promoting sparsity in the recovered solution. Priorconditioned (prior-preconditioned) hybrid methods have been proposed to improve performance, but introduce additional computational costs in each iteration related to the application of a weighted pseudoinverse to a matrix of basis vectors. We propose to use multigrid-preconditioned block Krylov methods to accelerate the application of the weighted pseudoinverses in each iteration, which are suitable

for both distributed CPU and GPU computing environments. Numerical results are presented for large-scale static and dynamic inverse problems.

Parametric Hyperbolic Conservation Laws: Learning Hyperbolic Conservation Laws from Data through Symmetrization

Authors: Anne Gelb[1], Lizuo Liu[1], Lu Zhang[2]

[1] *Dartmouth College*

[2] *Rice University*

Contact: lizuo.liu@dartmouth.edu

We propose a parametric hyperbolic conservation law (SymCLaw) for learning hyperbolic systems directly from data while ensuring conservation, entropy stability, and hyperbolicity by design. Unlike existing approaches that typically enforce only conservation or rely on prior knowledge of the governing equations, our method parameterizes the flux functions in a form that guarantees real eigenvalues and complete eigenvectors of the flux Jacobian, thereby preserving hyperbolicity. At the same time, we embed entropy-stable design principles by jointly learning a convex entropy function and its associated flux potential, ensuring entropy dissipation and the selection of physically admissible weak solutions. A corresponding entropy-stable numerical flux scheme provides compatibility with standard discretizations, allowing seamless integration into classical solvers. Numerical experiments on benchmark problems, including Burgers', shallow water, Euler, and KPP equations, demonstrate that SymCLaw generalizes to unseen initial conditions, maintains stability under noisy training data, and achieves accurate long-time predictions, highlighting its potential as a principled foundation for data-driven modeling of hyperbolic conservation laws.

Mixed-to-Low Precision Iterative Methods for Linear Inverse Problems

Authors: Lucas Onisk[1], Chelsea Drum[1], James Nagy[1]

[1] *Emory University*

Contact: lonisk@emory.edu

Many problems in science and engineering give rise to linear systems of equations that are commonly referred to as large-scale linear discrete ill-posed problems. The matrices that define these problems are typically severely ill-conditioned and may be rank deficient. Because of this, regularization is often needed to stem the effect of perturbations caused by error in the available data. In this talk we consider the solution of the regularized least-squares problem using both mixed-to-low precision and Krylov subspace projection techniques. We utilize a filter factor analysis to investigate the

regularizing behavior of the proposed iterative methods as well as numerical experiments to verify their efficacy.

Data-driven discovery of chemical reaction networks

Authors: Abraham Reyes Velazquez[1], Igor Larrosa[1], Jonas Latz[1], Stefan Guettel[1]

[1] *University of Manchester*

Contact: abrahamrafael.reyesvelazquez@manchester.ac.uk

We propose a unified framework that allows for the full mechanistic reconstruction of chemical reaction networks (CRNs) from concentration data. The framework utilizes an integral formulation of the differential equations governing the chemical reactions, followed by an automatic procedure to recover admissible mass-action mechanisms from the equations. We provide theoretical justification for the use of integral formulations using analytical and numerical error bounds. The integral formulation is demonstrated to offer superior robustness to noise and improved accuracy in both rate-law and graph recovery when compared to other commonly used formulations. Together, our developments advance the goal of fully automated, data-driven chemical mechanism discovery.

Bayesian inference for rough feature reconstructions

Author: Lassi Roininen[1]

[1] *LUT University*

Contact: lassi.roininen@lut.fi

Edges in imaging, that is sharp discontinuities in intensity, pose a significant challenge for inverse problems algorithms that often rely on Gaussian assumptions. Non-Gaussian heavy-tailed priors, which can better model the sparsity and sharp transitions inherent in edges, offer an alternative for edge-preserving image reconstructions. We consider the inherent difficulties in handling edges and highlight the potential of heavy-tailed prior models to transform this problem into a practical engineering solution.

Discretization-free Bayesian inverse problems

Authors: Erkki Somersalo[1], Daniela Calvetti[1]

[1] *Case Western Reserve University*

Contact: ejs49@case.edu

In this talk, we revisit the Bayesian inverse problems formalism in infinite-dimensional distribution spaces, where function evaluations are replaced by evaluations by test functions. It is shown that linear inverse problems can be formulated without a reference to any infinite-dimensional representation of the unknown, e.g., in terms of basis vectors, and therefore, the forward problem has a matrix-free form. In the Gaussian case, the numerical evaluation of the posterior mean and covariance matrix does not rely on finite-dimensional approximations of the unknown, but rather on numerical quadratures. The formalism is flexible, allowing a posteriori probing of the unknown without recalculation of any matrix inverses.

Nonlinear OED with Column Subset Selection

Authors: Amit Subrahmanya[1], Arvind Krishna Saibaba[2], Srinivas Eswar, Vishwas Rao

[1] *Argonne National Laboratory*

[2] *North Carolina State University*

Contact: famitnagesh@anl.gov

We address optimal sensor placement for Bayesian nonlinear inverse problems by formulating the task as a matrix column subset selection problem. The design matrix is derived from the expected information gain criterion. Although the resulting solutions are not necessarily globally optimal, the approach presents a rapid time to solution. The effectiveness of the method is demonstrated on nonlinear model problems.

Randomized Generalized Error Minimizing Method for Linear Ill-Posed Problem

Author: Ning Zheng[1]

[1] *Tongji University*

Contact: nzheng@tongji.edu.cn

For solving noisy linear ill–posed problems arising from the practical applications, the residual based iterative methods may suffer semi-convergence phenomenon, where the iterates initially get closer to the desired solution but then degrade as the iteration continues. Building upon the randomized Gram–Schmidt algorithm, a random sketching technique known to reduce inner product computational costs over classical Gram–Schmidt and numerical stability comparable to modified Gram–Schmidt, we develop a novel randomized generalized error minimizing (GMERR) Krylov subspace method. This process extends the successful application of randomized Gram–Schmidt in methods such as randomized GMRES and LSQR. We further introduce a block variant, resulting in

a block randomized Arnoldi process and a block GMERR method for large-scale ill-posed problems. A theoretical analysis of the regularization properties and numerical stability of the proposed methods is provided, leveraging random projection theory. Numerical experiments demonstrate the efficacy of the new algorithms.

MS 11: Numerical Linear Algebra Tools for Model Order Reduction

Minisymposium Organizers: Mattia Manucci, Sean Reiter

Modern scientific and engineering simulations increasingly rely on large systems of differential or algebraic equations derived from high-fidelity mathematical models. Although such models capture complex physical phenomena with great accuracy, their high dimensionality often renders them computationally infeasible for downstream tasks such as real-time control, optimization, or uncertainty quantification. Model-order reduction (MOR) provides a systematic framework for building low dimensional surrogate models that retain the essential input–to-output or dynamical behavior of the original systems at a fraction of the computational cost. The effectiveness and accuracy of most consolidated MOR routines are intricately connected to numerous numerical linear algebra methods, which play a vital role in the construction of low-dimensional surrogate models. This mini-symposium aims to bring together researchers who would showcase new and original integrations of numerical linear algebra routines in the development of low-dimensional surrogate models. In particular, we welcome and encourage participation from researchers who are in the early stages of their careers. The scope of interest includes the following areas: Krylov subspace techniques for model order reduction; recyclable singular value decomposition, which is essential for efficient proper orthogonal decomposition based MOR; novel advancements in matrix equation solvers for projection-based MOR; low-rank approximations and data compression; MOR for parametric eigenvalue problems; rational interpolation for data-driven surrogate modelling. By fostering dialogue and collaboration among early-stage researchers, this mini-symposium aims to identify synergistic opportunities between cutting-edge numerical linear algebra methods and MOR strategies to accelerate computational workflows without compromising accuracy.

List of Talks

Model Reduction For Optimal Control By Balanced Truncation Of State and Gradient Covariance

Authors: Emmanuel Ameh[1], Sam Otto[1]

[1] *Cornell University*

Contact: esa62@cornell.edu

Data-driven reduced-order models (ROMs) could enable near-optimal control for very high-dimensional nonlinear dynamical systems, with applications in active flow control such as relaminarizing turbulent flows and recovering from aerodynamic stall. With initial conditions far away from the desired steady state solving the resulting Hamilton-Jacobi-Bellman (HJB) equation, which defines the value function necessary for feedback control over the continuous state space, computationally intractable due to the curse of dimensionality. Reduced-order models (ROMs) can help in a variety of ways, but existing methods often fail to capture relevant dynamics for the control problem.

To overcome these challenges, we employ an indirect method of trajectory optimization (which is feasible in high dimensions) to obtain state and costate data offline, along locally optimal trajectories for estimating state and gradient covariance matrices. This method builds on the Pontryagin minimum principle and other related work that establishes the costate (adjoint variables) provided as generalized gradients of the optimal value function satisfying the HJB equation. An oblique projection obtained by balanced truncation of these matrices with initial conditions sampled uniformly from an uncontrolled attractor, is used to identify active directions in the state space along which the value function is most sensitive and states have large variance.

The oblique projection obtained is used to build surrogate models for both the value function and the optimal feedback control law, which is validated on the full order model. ROM-based state estimators are also built with these projections for closed-loop feedback control. We assess the quality of the resulting ROMs across linear and nonlinear flow control problems, benchmarking against existing ROM approaches.

On linear matrix equations and the Akhiezer iteration

Authors: Cade Ballew[1], Heather Wilber[1], Thomas Trogdon[1]

[1] *University of Washington*

Contact: ballew@uw.edu

The Akhiezer iteration is a new iterative method for solving indefinite linear systems and computing matrix functions. The iteration uses orthogonal polynomial recurrence coefficients to efficiently

compute the action of a matrix polynomial to a vector without computing inner products. It features an a priori computable convergence rate and is often faster in practice than standard Krylov subspace methods.

We present extensions of the Akhiezer iteration to large-scale linear matrix equations, such as the Sylvester equation $AX + XB = C$, arising from model order reduction problems.

PIRKA: The Iterative Rational Krylov Algorithm for Linear Time-Periodic Systems

Authors: Christopher Beattie, Sam Bender
Contact: skbender@vt.edu

We consider single-input, single-output systems with time-varying, periodic parameters:

$$\dot{x}(t) = A(t)x(t) + b(t)u(t),$$
$$y(t) = c(t)x(t),$$

where $A(t) \in \mathbb{R}^{n\times n}$ and $b(t), c(t) \in \mathbb{R}^n$ all have period T. Such systems arise when modeling phenomena in fluid dynamics, structural mechanics, and electronic circuits. In particular, linearization around known periodic orbits of a nonlinear model produces a periodic system of partial differential equations, and subsequent spatial semi-discretization yields large-scale linear time-periodic (LTP) dynamical systems. The need to simulate system responses to a variety of inputs motivates the development of effective model reduction tools for these systems. While research on model reduction for LTP systems is relatively limited, there is a substantial body of literature devoted to control, spectral analysis, and harmonic response of LTP systems. From this literature arise the Harmonic Transfer Function and the H_2 norm for LTP systems. These concepts lead to necessary conditions for an H_2-optimal reduced-order model. In the linear time-invariant case, the Iterative Rational Krylov Algorithm (IRKA) is a standard approach for the H_2 model reduction problem. The work presented here extends IRKA to the LTP setting.

Data-Driven Modal Truncation

Authors: Dan E. Folescu[1], Mark Embree[1], Michael C. Brennan[2], Serkan Gugercin[1]

[1] *Virginia Tech*

[2] *Solea Energy*

Contact: dan1345@vt.edu

Modal truncation has long been a fundamental approach of model order reduction: to systematically eliminate eigenmodes of a dynamical system that contribute little to the modeling behavior over

a given time/frequency range. Typically, this procedure requires the ability to access and utilize *intrusive* state-space information about the underlying full-order system, which can be infeasible for larger state dimensions or impossible if the full-order system is not known in closed form. Data-driven approaches like proper orthogonal decomposition address these concerns by identifying *empirical* modes of the system derived from multiple time-domain trajectories of the full-order model, whereafter model reduction can proceed as in the state-space formulation. In this work, we show how one can truncate empirical eigenmodes with only frequency-domain samples of a system's transfer function by utilizing contour integral methods.

After reviewing the necessary theoretical foundations, we will discuss how contour integration and the pole-residue decomposition of an underlying system can be combined to identify all eigenvalues within an arbitrary domain in the complex plane. This tool allows us to perform modal truncation, or identify a stable-unstable decomposition of the underlying system. Through the use of a key result of Keldysh, contour integral methods can also identify eigenvalues of analytic nonlinear matrix-valued functions. We will show numerical experiments illustrating how our data-driven modal truncation method performs on application problems and compare with other state-of-the-art methods.

Efficient Streaming Operator Learning for Large-Scale Dynamical Systems

Authors: Tomoki Koike[1], Prakash Mohan[2], Marc T. Henry de Frahan[2], Julie Bessac[2], Elizabeth Qian[1]

[1] *Georgia Institute of Technology*

[2] *National Laboratory of the Rockies*

Contact: tkoike3@gatech.edu

Modeling and simulation of real-world applications often involve dynamical systems with large degrees of freedom, requiring substantial computational time and resources. Projection-based model reduction enables efficient simulation of such dynamical systems by constructing low-dimensional surrogate models from high-dimensional data. Specifically, Operator Inference (OpInf) learns such reduced surrogate models through a two-step process: constructing a low-dimensional basis via Singular Value Decomposition (SVD) to compress the data, then solving a linear least-squares (LS) problem to infer reduced operators that govern the dynamics in this compressed space, all without access to the underlying code or full model operators, i.e., non-intrusively. Traditional OpInf operates as a batch learning method, where both the SVD and LS steps process all data simultaneously, which limits scalability to large-scale applications generating terabytes to petabytes of data and prevents real-time model updates in online scenarios. To address these limitations, we propose Streaming OpInf, which learns reduced models incrementally as snapshot data arrives. Our method employs prominent streaming algorithms from numerical linear algebra: incremental SVD for adaptive basis construction and recursive LS for streaming operator updates, eliminating the need to store complete datasets while enabling online model adaptation. We systematically compare multiple streaming algorithm variants to identify effective combinations for accurate reduced

model learning. A Numerical experiment on a large-scale turbulent channel flow with a friction Reynolds number of $Re_\tau = 5200$ demonstrate that Streaming OpInf achieves accuracy comparable to batch OpInf while reducing memory requirements significantly and enabling dimension reductions exceeding 31,000x, resulting in orders-of-magnitude faster predictions.

Solving Generalized Lyapunov Equations with guarantees: application to the Model Reduction of Switched Linear Systems

Authors: Mattia Manucci[1], Benjamin Unger[1]

[1] *Karlsruhe Institute of Technology*

Contact: mattia.manucci@simtech.uni-stuttgart.de

In this talk, we present an efficient strategy to approximate the solutions of large-scale generalized Lyapunov equations (GLEs) while providing rigorous error guarantees. The motivation for this study stems from the use of GLEs in model order reduction (MOR) of switched linear systems (SLS) in control form. Specifically, we analyze how inaccuracies in the computed GLE solution influence the performance and reliability of the resulting MOR procedure. Furthermore, since the classical balanced-truncation error estimate for MOR of SLS is neither theoretically nor practically viable—because it relies on restrictive assumptions that require several linear matrix inequalities (LMI) to be satisfied by numerically computed solutions of the GLEs—we propose a new MOR framework for SLS. Our method is based on solving multiple GLEs and constructing projection matrices that are piecewise constant in time to appropriately balance and subsequently reduce the SLS; we therefore refer to it as piecewise balancing reduction (PBR) for SLS. We extend the standard balance-truncation error bounds to incorporate the effects of inexact LMI. We show how the PBR formulation allows us to control the error arising from the inexact LMI. In addition, our new error bound accounts for the influence of the piecewise constant time-varying projection matrices. Altogether, this renders the PBR approach for SLS applicable to a broad and flexible class of SLS. We conclude by showing numerical experiments to corroborate our theoretical results.

Least-squares Rational Approximation Using Riemannian Optimization

Author: Petar Mlinarić[1]

[1] *University of Zagreb*

Contact: petar.mlinaric@math.hr

Vector fitting is a widely used method for least-squares rational approximation that approaches the nonlinear least-squares problem with a sequence of linear least-squares problems. By contrast, the iterative rational Krylov algorithm (IRKA) is a method for a continuous least-squares

approximation, originally formulated as a fixed-point iteration and with a recent interpretation as a Riemannian optimization method. In this talk, we discuss the application of Riemannian optimization techniques to the (discrete) least-squares rational approximation problem and compare them with vector fitting.

Extensions of data-driven balancing to LQO and QB systems

Authors: Reetish Padhi[1], Ion Victor Gosea[2], Igor Pontes Duff[3], Serkan Gugercin[1]

[1] *Virginia Tech*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany*

[3] *CSC group, Max Planck Institute of Dynamics of Complex Technical Systems, Magdeburg, Germany*

Contact: reetishp@vt.edu

We develop the theoretical framework for extending the quadrature based balanced truncation (QuadBT) method to linear systems with quadratic outputs (LQO). QuadBT which was originally designed for data-driven balanced truncation of standard linear systems with linear outputs only. We show that by sampling the extended impulse responses (kernels) and their derivatives (in the time domain) or the corresponding transfer functions (in the frequency domain), we can construct a reduced order model that mimics the approximation quality of intrusive balanced truncation. At its core, the method can be interpreted as an implicit sampling of the Gramians of the system using input-output observations. We demonstrate a proof of concept for the proposed framework on an example using numerically evaluated data. We also briefly discuss the extension of the QuadBT method for the more involved quadratic-bilinear (QB) system case.

$\mathcal{H}_2$ model order reduction for Bilinear Quadratic Output Systems

Authors: Till Peters[1], Heike Faßbender, Serkan Gugercin[2]

[1] *TU Braunschweig, Institute for Numerical Analysis*

[2] *Virginia Tech*

Contact: till.peters@tu-braunschweig.de

Today, mathematical modeling is dominated by increasingly high-dimensional and complex dynamical systems. One special type of structure is the bilinear state equation, which either naturally appears in various applications or results from the Carleman bilinearization of the underlying nonlinear dynamics. Recently, dynamical systems with quadratic outputs have also gained significant attention as they appear, e.g., in the modeling of the variance of a quantity of interest of a stochastic model. As a combination, we study the so-called bilinear quadratic output (BQO) systems described as

$$\dot{x}(t) = Ax(t) + Nx(t)u(t) + Bu(t), y(t) = Cx(t) + x(t)^T Mx(t)$$

with state $x(t) \in \mathbb{R}^n$, scalar input $u(t)$ and output $y(t)$ and matrices A, B, C, N, and M of suitable dimensions (SISO case).

BQO systems generalize bilinear systems ($M = 0$) and linear quadratic output systems ($N = 0$). For large dimensional systems, the task of structure-preserving model order reduction arises. Thus, we seek a BQO system of reduced dimension that approximates the original input-output behavior well. Whereas a balancing approach has been already discussed in [2], here we focus on an $\mathcal{H}_2$ optimal approach, similar to $\mathcal{H}_2$ optimal reduction for bilinear systems; see, e.g., [1], [3]. We will establish the framework for the $\mathcal{H}_2$ model order reduction of BQO systems, develop iterative algorithms for finding a locally optimal reduced system, and test these algorithms on several numerical examples.

References:
[1] P. Benner and T. Breiten. Interpolation-based $\mathcal{H}_2$-model reduction of bilinear control systems. SIAM J. Matrix Anal. Appl., 33(3):859-885, 2012.
[2] H. Faßbender, S. Gugercin, and T. Peters. Bilinear Quadratic Output Systems and Balanced Truncation, July 2025. arXiv:2507.03684 [math].
[3] L. Zhang and J. Lam. On $\mathcal{H}_2$ model reduction of bilinear systems. Automatica, 38(2):205-216, Feb. 2002.

A tangential low-rank ADI method for solving indefinite Lyapunov equations

Authors: Rudi Smith[1], Steffen W. R. Werner[1]

[1] *Virginia Tech*

Contact: smithrgh@vt.edu

Continuous-time algebraic Lyapunov equations are linear matrix equations of the form

$$AXE^{\mathsf{H}} + EXA^{\mathsf{H}} = -W$$

where $A, E \in \mathbb{C}^{n\times n}$ are large-scale sparse coefficient matrices and $W = BRB^{\mathsf{H}}$ represents an indefinite right-hand side defined by the low-rank factor $B \in \mathbb{C}^{n\times m}$ and the indefinite Hermitian matrix $R = R^{\mathsf{H}} \in \mathbb{C}^{m\times m}$. Such matrix equations have become an essential tool in various applications like model reduction, control or data-driven modeling. In the case of large-scale sparse coefficient matrices and indefinite constant terms, indefinite low-rank factorizations have successfully been used to allow methods like the alternating direction implicit (ADI) iteration to efficiently compute accurate approximations to the solution of the Lyapunov equation. However, classical block-type approaches quickly increase in computational costs when the rank of the constant term grows. In this work, we propose a novel tangential reformulation of the ADI iteration that allows for the efficient construction of low-rank approximations to the solution of Lyapunov equations with indefinite right-hand sides even in the case of constant terms with higher ranks. We provide adaptive

methods for the selection of the corresponding ADI parameters, namely shifts and tangential directions, which allow for the automatic application of the method to any relevant problem setting. The effectiveness of the developed algorithms is illustrated by several numerical examples.

Inverse problems for history-enriched linear model reduction

Authors: Arjun Vijaywargiya[1], George Biros[1]

[1] *University of Texas at Austin*

Contact: arjunveejay@utexas.edu

Standard projection-based model reduction for dynamical systems incurs closure error because it only accounts for instantaneous dependence on the resolved state. From the Mori–Zwanzig (MZ) perspective, projecting the full dynamics onto a low-dimensional resolved subspace induces additional noise and memory terms arising from the dynamics of the unresolved component in the orthogonal complement. The memory term makes the resolved dynamics explicitly history dependent. In this work, based on the MZ identity, we derive exact, history-enriched models for the resolved dynamics of linear driven dynamical systems and formulate inverse problems to learn model operators from discrete snapshot data via least-squares regression. We propose a greedy time-marching scheme to solve the inverse problems efficiently and analyze operator identifiability under full and partial observation data availability. For full observation data, we show that, under mild assumptions, the operators are identifiable even when the full-state dynamics are governed by a general time-varying linear operator, whereas with partial observation data the inverse problem has a unique solution only when the full-state operator is time-invariant. To address the resulting non-uniqueness in the time-varying case, we introduce a time-smoothing Tikhonov regularization. Numerical results demonstrate that the operators can be faithfully reconstructed from both full and partial observation data and that the learned history-enriched MZ models yield accurate trajectories of the resolved state.

MS 12: Theoretical Advances in Operator Learning

Minisymposium Organizers: Diana Halikias, Sam Otto

The rapidly developing field of operator learning addresses the following question: what can be learned about an operator that maps between infinite-dimensional spaces, such as a dynamical system or the solution operator of a partial differential equation, from only observational data? Resolving this question allows one to build efficient surrogate models and accelerated solvers for complex phenomena that are challenging to model with traditional methods. When the underlying operator is truly unknown, operator learning may even uncover unknown physical laws and quantities. This minisymposium is focused on recent developments surrounding the theoretical underpinnings of operator learning, which have important connections to linear algebra. Talks will feature ideas from sketching/randomized linear algebra, approximation theory, optimization, and statistical learning in high/infinite dimensions, with applications to inverse problems and dynamical systems.

List of Talks

Operator Learning at Machine Precision

Authors: Aras Bacho[1], Aleksei Sorokin, Bamdad Hosseini, Edoardo Calvello, Matthieu Darcy, Theo Bourdais, Xianjin Yang, Houman Owhadi, Alexander Hsu

[1] *California Institute of Technology*

Contact: bacho@caltech.edu

Neural operator learning methods have garnered significant attention in scientific computing for their ability to approximate infinite-dimensional operators. However, increasing their complexity often fails to substantially improve their accuracy, leaving them on par with much simpler approaches such as kernel methods and more traditional reduced-order models. In this article, we set out to address this shortcoming and introduce CHONKNORIS (Cholesky Newton–Kantorovich Neural Operator Residual Iterative System), an operator learning paradigm that can achieve machine precision. CHONKNORIS draws on numerical analysis: many nonlinear forward and inverse PDE problems are solvable by Newton-type methods. Rather than regressing the solution operator itself, our method regresses the Cholesky factors of the elliptic operator associated with Tikhonov-regularized Newton–Kantorovich updates. The resulting unrolled iteration yields a neural architecture whose machine-precision behavior follows from achieving a contractive map, requiring far lower accuracy than end-to-end approximation of the solution operator. We benchmark CHONKNORIS on a range of nonlinear forward and inverse problems, including a nonlinear elliptic equation, Burgers' equation, a nonlinear Darcy flow problem, the Calder\'{o}n problem, an inverse wave scattering problem, and a problem from seismic imaging. We also present theoretical guarantees for the convergence of CHONKNORIS in terms of the accuracy of the emulated Cholesky factors. Additionally, we introduce a foundation model variant, FONKNORIS (Foundation Newton–Kantorovich Neural Operator Residual Iterative System), which aggregates multiple pre-trained CHONKNORIS experts for diverse PDEs to emulate the solution map of a novel nonlinear PDE. Our FONKNORIS model is able to accurately solve unseen nonlinear PDEs such as the Klein–Gordon and Sine–Gordon equations.

Estimation for intrinsic Gaussian processes

Authors: Christopher Beattie[1], David Higdon[1], Leanna House[1]

[1] *Virginia Tech*

Contact: beattie@vt.edu

Gaussian processes (GPs) defined through intrinsic random fields provide a flexible framework for modeling spatial phenomena, and have been advocated in a variety of applications over the past several decades. Nevertheless, their adoption has lagged behind traditional, covariance-based approaches, in part because the intrinsic formulation has lacked an accompanying toolkit of computational methods and dependence specifications that facilitate fitting and prediction. This work develops a systematic framework for modeling intrinsic GPs and introduces practical algorithms for working with dependence/variogram models for modeling, inference and computation that parallel those of traditional, stationary GPs, highlighting the advantages of intrinsic-field modeling in terms of robustness, interpretability, and computational efficiency.

Resolvent compactification methods for spectral approximation of Koopman operators

Author: Dimitrios Giannakis[1], Trevor Camper[1]

[1] *Dartmouth College*

Contact: trevor.camper@dartmouth.edu

Koopman operators and transfer operators represent nonlinear dynamics in state space through its induced action on linear spaces of observables and measures, respectively. This framework enables the use of linear operator theory for analysis and modeling of nonlinear dynamical systems, and has received considerable interest over the years from mathematical, computational, and domain-scientific disciplines. In this talk, we present data-driven techniques for spectral approximation of Koopman and transfer operators of measure-preserving ergodic flows that are based on operators with compact resolvent. Our approach performs a bounded transformation of the Koopman generator (an operator implementing directional derivatives of observables along the dynamical flow), followed by smoothing by a Markov semigroup of kernel integral operators. This results in a skew-adjoint, compact operator with compact resolvent whose eigendecomposition is expressible as a variational generalized eigenvalue problem amenable to approximation using Galerkin methods. A key aspect of these methods is that they are physics-informed, in the sense of making direct use of dynamical vector field information through automatic differentiation of kernel functions. Solutions of the eigenvalue problem reconstruct evolution operators that preserve unitarity of the underlying Koopman group while spectrally converging to it in a suitable limit. In addition, the computed eigenfunctions have representatives in a reproducing kernel Hilbert space, enabling out-of-sample evaluation of learned dynamical features. We illustrate this method with numerical experiments on integrable and chaotic systems.

Learning Enhanced Ensemble Filters: Continuum Limits of Attention on Measure

Authors: Bohan Chen[1], Andrew Stuart[1], Edoardo Calvello[1], Eviatar Bach[2], Ricardo Baptista[3]

[1] *California Institute of Technology*

[2] *University of Reading*

[3] *University of Toronto*

Contact: bhchen@caltech.edu

Many classic methods in data assimilation, like the Ensemble Kalman Filter (EnKF), are limited by its Gaussian ansatz. In this work, we frame the filtering update as learning a nonlinear operator mapping between probability distributions in the mean-field limit. We introduce Measure Neural Mappings (MNMs), a class of neural operators acting on probability measures, implemented via Set Transformers. A central theoretical contribution is establishing the continuum limit for attention mechanisms on measures. We prove that attention layers applied to finite ensembles are consistent

with their continuous-measure counterparts, converging in Wasserstein distance as the sample size approaches infinity. This result rigorously links finite-dimensional linear algebra operations to infinite-dimensional operator learning, justifying the use of a single parameterization across different ensemble sizes. We demonstrate that the resulting MNM-enhanced ensemble filter (MNMEF) achieves superior accuracy on chaotic dynamical systems, including the Lorenz '96 and Kuramoto-Sivashinsky models, outperforming leading filtering methods.

Data-efficient Adjoint-free Learning for Asymptotically Smooth Integral Operators

Authors: Esther Gallmeier[1], Samuel Otto[1]

[1] *Cornell University*

Contact: eg542@cornell.edu

The goal underlying our work is to develop provably accurate and data-efficient learning algorithms for non-self-adjoint operators using only input-output pairs. State-of-the-art approximation techniques with fast convergence rates either apply only to self-adjoint operators or require access to the adjoint operator, which is unavailable in experimental settings and often difficult to access computationally. Recent work has shown that provably accurate adjoint-free learning is possible but uses an amount of data that is likely excessive. In this talk, we discuss an adjoint-free learning algorithm for asymptotically smooth integral operators which is based on an extension of the fixed sparsity matrix approximation algorithm [Amsel et al., arXiv, 2024] to the case where a fixed subspace - determined by asymptotic smoothness and the resulting exponentially-convergent separable expansion - is known for each row. The algorithm queries the action of the underlying operator on functions with standard Gaussian random coefficients in a finely discretized finite element space, yielding an approximation of the unknown finite element discretization of the operator. While the number of degrees of freedom in this discretization grows polynomially with the inverse error tolerance, asymptotic smoothness and the hierarchical decomposition of each row ensure that only a polylogarithmic number of queries $\mathcal{O}(\text{polylog}(\varepsilon^{-1}))$ is required to achieve the desired accuracy ε with high probability.

Towards High-Precision Optimizers for Scientific Machine Learning

Authors: Gil Goldshlager[1], Jiang Hu[2], Lin Lin[1]

[1] *UC Berkeley*

[2] *Tsinghua University*

Contact: ggoldshlager@gmail.com

An increasing number of theoretical results are available to characterize the extent to which neural networks can (i) represent scientifically relevant functions and operators, and (ii) learn these functions and operators from data. However, even with the right network architecture and the right dataset, optimization is a bottleneck. On the one hand, popular machine learning optimizers such as SGD and Adam are not designed to attain the levels of precision needed for scientific applications. On the other hand, traditional approaches to high-precision optimization, such as Newton or quasi-Newton methods, are not designed for the highly stochastic training regimes that are characteristic of machine learning. In this talk, I will argue that these limitations can be overcome by adopting randomized linear algebra as a paradigm for high-precision scientific machine learning. As an example, I will share some recent work which uses this paradigm to both explain and improve a promising class of optimizers known as subsampled natural gradient algorithms.

L2-Optimal Reduced-Order Modeling Using Parameter-Separable Forms

Authors: Serkan Gugercin[1], Petar Mlinarić[2]

[1] *Virginia Tech*

[2] *University of Zagreb*

Contact: gugercin@vt.edu

We introduce a unified approach to $\mathcal{L}_2$-optimal reduced-order modeling that applies to both linear time-invariant dynamical systems and stationary parametric problems. The framework leverages parameter-separable representations to obtain gradient information for the $\mathcal{L}_2$ objective with respect to the reduced operators, enabling a fully nonintrusive, data-driven, gradient-based construction of optimal reduced models from output data alone. By selecting an appropriate measure, the formulation naturally includes both continuous and discrete cost functions. The proposed methodology is validated through representative numerical examples, and conditions guaranteeing a projection-based realization of the data-driven approximant are established.

Learning operators with continuous spectrum from data

Author: Andrew Horning[1]

[1] *Rensselaer Polytechnic Institute*

Contact: hornia3@rpi.edu

Linear operators with a continuous spectrum often lurk behind complex physical phenomena in nature, from wave attractors in geometrically confined fluids to topological bifurcations in dynamical systems. However, they are notoriously tricky to learn from data. For example, finite-dimensional

approximations of the operator must "discretize" the continuous spectrum into finitely many points and cannot converge uniformly.

In this talk, we explore two principled approaches to data-driven approximation of unitary operators with continuous spectrum. The first approach synthesizes recent advances in computational spectral theory with stable algorithms for Gauss-Szego quadrature rules on the unit circle. The second approach leverages new algorithms for quadrature rules that adaptively deform and discretize the contour inside the unit circle, resembling classical resonance expansions for wave operators. We discuss strong convergence rates for the approximate operators, data-driven aspects of the approximation, and highlight the unifying power of quadrature rules derived from rational approximations.

Operator Learning via Learned Differential Operators

Authors: Alexander Hsu[1], Bamdad Hosseini[1], Houman Owhadi[2], Juan Felipe Osorio Ramirez[1], Yasamin Jalalian[2]

[1] *University of Washington*

[2] *Caltech*

Contact: josorior@uw.edu

We present an alternative perspective on operator learning for problems in which the operator is implicitly defined by a partial differential equation. Rather than learning the solution operator directly as a high-dimensional mapping, we propose to first learn the underlying PDE operator as a local differential operator and then numerically invert it to evaluate the associated solution operator.

Quasi-optimal hierarchically semi-separable matrix approximation

Authors: Noah Amsel[1], Tyler Chen[2], Feyza Duman Keles[1], Diana Halikias[1], Cameron Musco[3], Christopher Musco[1], David Persson[4]

[1] *New York University*

[2] *JP Morgan Chase*

[3] *University of Massachusetts Amherst*

[4] *New York University & Flatiron Institute*

Contact: dup210@nyu.edu

We present a randomized algorithm for producing a quasi-optimal hierarchically semi-separable (HSS) approximation to an $N \times N$ matrix A using only matrix-vector products with A and A^T. We prove that, using $O(k \log(N/k))$ matrix-vector products and $O(Nk^2 \log(N/k))$ additional runtime,

the algorithm returns an HSS matrix B with rank-k blocks whose expected Frobenius norm error $\mathbb{E}[\|A - B\|_F^2]$ is at most $O(\log(N/k))$ times worse than the best possible approximation error by an HSS rank-k matrix. In fact, the algorithm we analyze in a simple modification of an empirically effective method proposed by [Levitt & Martinsson, SISC 2024]. As a stepping stone towards our main result, we prove two results that are of independent interest: a similar guarantee for a variant of the algorithm which accesses A's entries directly, and explicit error bounds for near-optimal subspace approximation using *projection-cost-preserving sketches*. To the best of our knowledge, our analysis constitutes the first polynomial-time quasi-optimality result for HSS matrix approximation, both in the explicit access model and the matrix-vector product query model.

Learning Material Constitutive Laws with Neural Operators

Author: George Stepaniants[1]

[1] *California Institute of Technology*

Contact: gstepan@caltech.edu

The simulation of multiscale viscoelastic materials poses a significant challenge in computational materials science, requiring expensive numerical solvers that can resolve dynamics of material deformations at the microscopic scale. The theory of homogenization offers an alternative approach to modeling, by locally averaging the strains and stresses of multiscale materials. This procedure eliminates the smaller scale dynamics but introduces a history dependence between strain and stress that proves very challenging to characterize analytically. In the one-dimensional setting, we give the first full characterization of the memory-dependent constitutive laws that arise in multiscale viscoelastic materials. Using this theory, we develop a neural differential equation architecture, that simultaneously across a wide range of material microstructures, accurately predicts their homogenized constitutive laws, thus enabling us to simulate their deformations under forcing. We use the approximation theory of neural operators to provide guarantees on the generalization of our approach to unseen material samples.

MS 13: Computational Advances in Discrete Inverse Problems

Minisymposium Organizers: Silvia Gazzola, Lucas Onisk, Malena Sabate Landman

Inverse problems remain central to scientific modeling and data interpretation, yet their inherent ill-posedness and large-scale structure continue to drive the need for innovative computational strategies. Recent advances in randomized algorithms, mixed-precision computation, Bayesian inference, and scalable numerical methods are reshaping how discrete inverse problems are formulated and solved. This minisymposium brings together researchers developing cutting-edge techniques that balance computational efficiency with accuracy and stability in challenging, high-dimensional settings. By highlighting both emerging methodologies and their practical impacts, the session aims to foster cross-disciplinary dialogue and inspire new directions in the broader inverse problems' community.

List of Talks

Latent Twin Operator

Authors: Deepanshu Verma[1], Joseph Lee Hart[2], Matthias Chung[3], Riley Yizhou Chen[3]

[1] *Clemson University*

[2] *Sandia National Laboratories*

[3] *Emory University*

Contact: rileychen111@gmail.com

Learning solution operators in a manner that is independent of discretization and resolution remains a central challenge in data-driven modeling. The latent twins framework addresses this problem by constructing operators in a task-adaptive latent space for inverse problems and differential

equations. However, in its classical form, latent twins relies on autoencoder architectures that are tied to fixed discretizations, coupling representation learning to a particular grid or resolution.

We address this limitation by separating state information from coordinate information. The latent variables encode the global system state, while a coordinate-conditioned decoder acts as an evaluation operator that reconstructs the state at arbitrary spatial or temporal locations through a family of linear or nonlinear maps. This viewpoint naturally supports sparse, irregular, and multi-resolution data, and connects latent twins to operator learning and reduced-order modeling perspectives. The resulting framework is well suited for applications in imaging, inverse problems, and the reduced-order modeling of time-dependent systems.

Kronecker Rank Bounds for Operator Matrices

Authors: Allison Fuller[1], Malena Espanol[1], Misha Kilmer[2]

[1] *Arizona State University*

[2] *Tufts University*

Contact: malena.espanol@asu.edu

Block-structured matrices arise as operators in many contexts, including image deblurring and discretized differential equations. These matrices are often large and computationally difficult to work with. By rewriting these operators as a sum of Kronecker products, we may be able to alleviate these challenges. In this talk, we show how we can use the structure of a matrix to impose bounds on such a sum.

A Regularization Method for Compact Linear Operator Equations Based on the Arnoldi Process

Authors: Lothar Reichel[1], Mykhailo Kuian[2], Ronny Ramlau[3]

[1] *Kent State University*

[2] *Case Western Reserve University*

[3] *Industrial Mathematics Institute, Johannes Kepler University, and Johann Radon Institute for Computational and Applied Mathematics, Austrian Academy of Sciences*

Contact: michaelkuian@gmail.com

We consider the numerical solution of linear operator equations involving compact operators. Since compact operators do not admit bounded inverses, the associated equations are ill-posed and require regularization. The Arnoldi process provides a natural framework for approximating a compact operator by a nearby operator of finite rank, thereby reducing the infinite-dimensional problem to a

sequence of small, structured subproblems. Regularization is incorporated by applying Tikhonov's method to the projected equations.

This work investigates theoretical properties of the resulting Arnoldi–Tikhonov approach, including convergence behavior and the influence of Krylov subspace dimension on the regularized solutions. Numerical experiments are presented to illustrate the theoretical results and to demonstrate the effectiveness of the method for representative compact operator equations.

Regularized Krylov Subspace Methods in Low Precision

Authors: Ariana Brown[1], Chelsea Drum[1], James Nagy[1], Lucas Onisk[1]

[1] *Emory University*

Contact: jnagy@emory.edu

In recent years a substantial amount of work has been done on developing mixed-precision algorithms for linear systems, methods that can exploit capabilities of modern GPU architectures. However, very little work has been done for ill-conditioned problems that arise from large-scale inverse problems. Special considerations, which normally do not arise when solving well-conditioned problems, such as incorporating regularization into the developed methods, need to be considered. In this talk we consider the use of low-precision, regularized preconditioners in Krylov subspace methods.

Recycling and Streaming for Large Scale Nonlinear Inverse Problems

Authors: Eric de Sturler[1], James Nagy[2], Mirjeta Pasha, Misha Kilmer[3], Toluwani Okunola[3]

[1] *Virginia Tech*

[2] *Emory University*

[3] *Tufts University*

Contact: toluwani.okunola@tufts.edu

Many imaging inverse problems assume a known linear forward operator, yet practical systems often suffer from uncertainty in acquisition geometry, such as projection angles in computed tomography, sensor positions in photoacoustic tomography. These uncertainties introduce nonlinearity and require joint estimation of both the image and the forward model parameters.

We propose a nonlinear recycled majorization–minimization generalized Krylov subspace (NL-RMM-GKS) framework for large-scale inverse problems with uncertain forward operators. The method extends MM-GKS to nonlinear settings by combining majorization–minimization for non-smooth regularization with Krylov subspace projection and subspace recycling, ensuring bounded memory usage.

Two complementary formulations are developed: an alternating minimization approach that alternates between image updates and Gauss–Newton parameter estimation, and a variable projection approach that eliminates the image variable and optimizes directly over the parameters using inexact inner solves. We further introduce streaming variants that process data sequentially, enabling reconstruction from large or dynamically acquired datasets without storing the full operator.

We carry out rigorous numerical experiments in fan-beam computed tomography and photoacoustic tomography to demonstrate that our proposed framework achieves high-quality reconstructions with bounded memory requirements, making it suitable for large-scale and dynamic imaging problems with uncertain geometry.

IterativeCUR: One Small Sketch for Big Matrix Approximations

Authors: Nathaniel Pritchard[1], Per-Gunnar Martinsson[2], Taejun Park[3], Yuji Nakatsukasa

[1] *The University of Oxford*

[2] *University of Texas*

[3] *EPFL-Lausanne*

Contact: pritchard@maths.ox.ac.uk

The computation of accurate low-rank matrix approximations is central to improving the scalability of various techniques in machine learning, uncertainty quantification, and control. Traditionally, low-rank approximations are constructed using SVD-based approaches such as truncated SVD or RandomizedSVD. Although these SVD approaches—especially RandomizedSVD—have proven to be very computationally efficient, other low-rank approximation methods can offer even greater performance. One such approach is the CUR decomposition, which forms a low-rank approximation using direct row and column subsets of a matrix. Because CUR uses direct matrix subsets, it is also often better able to preserve native matrix structures like sparsity or non-negativity than SVD-based approaches and can facilitate data interpretation in many contexts. This paper introduces IterativeCUR, which draws on previous work in randomized numerical linear algebra to build a new algorithm that is highly competitive compared to prior work: (1) It is adaptive in the sense that it takes as an input parameter the desired tolerance, rather than an a priori guess of the numerical rank. (2) It typically runs significantly faster than both existing CUR algorithms and techniques such as RandomizedSVD, in particular when these methods are run in an adaptive rank mode. Its asymptotic complexity is $\mathcal{O}(mn + (m+n)r^2 + r^3)$ for an $m \times n$ matrix of numerical rank r. (3) It relies on a single small sketch from the matrix that is successively downdated as the algorithm proceeds.

New flexible and inexact Krylov solvers for inverse problems

Authors: Malena Sabate Landman[1], Silvia Gazzola[2]

[1] *University of Bath*

[2] *Università di Pisa*

Contact: msl39@bath.ac.uk

This talk presents a new family of algorithms for large-scale linear inverse problems built on flexible and inexact variants of the Golub–Kahan factorization. The proposed approach constructs regularized solutions through a sequence of projected (re)weighted least-squares problems, where the projection spaces are adaptively generated and endowed with iteration-dependent preconditioning and controlled inexactness. This framework enables a unified and flexible treatment of challenging problem settings, including general data fidelity models such as those involving p-norms. Numerical experiments in imaging applications, such as deblurring and computed tomography, highlight the effectiveness and competitiveness of the proposed methods with respect other popular methods.

MS 14: Code-based Cryptography

Minisymposium Organizers: Sarah Arpin, Jason LeGrow

Code-based cryptography is a promising new direction in post-quantum cryptography, with security rooted in hard linear algebra problems over finite fields. This session will feature research at the interface of cryptography, coding theory, and computational linear algebra to highlight recent advances, new techniques, and open questions. In 2025, the National Institute of Standards and Technology selected the code-based key encapsulation mechanism HQC for standardization. Previous code-based KEMs fell under the umbrella of either McEliece or Niederreiter-type cryptosystems, but HQC is a new type of KEM which combines three classes of linear codes (Reed-Solomon, Reed-Muller, and QC-MDPC codes). This new construction motivates the study of a broad spectrum of linear-algebraic properties: analyzing the spectral properties of the structured matrices, the decoding algorithms for concatenated codes, characterizing low-weight codewords and trapping sets, and exploiting structure to enable efficient algorithms. Our session will feature speakers working on these cutting-edge questions. The first practical code-based digital signatures were developed only within the past decade, making this a rapidly evolving research area. Recent schemes are built on variants of syndrome decoding, rank-metric problems, structured parity-check matrices, and code equivalence problems. These protocols rely heavily on linear-algebraic operations over finite fields. As these signature proposals mature, they highlight open questions and invite deeper structural analysis. Our session will feature researchers driving these new proposals. Our session's speakers will discuss advances in decoding algorithms, families of codes suitable for code-based cryptography, and cryptanalytic techniques that leverage the linear algebra structure underlying these codes over finite fields.

List of Talks

Decoding from the Other Side: Primal vs. Dual Attacks

Authors: Kévin Carrier[1], Charles Meyer-Hilfiger[2], Jean-Pierre Tillich[3], Laura Luzzi[4], Thomas Debris-Alazard[5], Valérian Hatey[6]

[1] *École Polytechnique (France)*

[2] *Inria Rennes*

[3] *Inria de Paris*

[4] *ENSEA Cergy*

[5] *Inria Saclay*

[6] *CY Cergy-Paris Université*

Contact: kevin.carrier@polytechnique.edu

The security of code-based cryptography relies fundamentally on the computational hardness of decoding random linear codes. Until recently, the most efficient known algorithms for the decoding problem were Information Set Decoding (ISD) algorithms, which we refer to as primal attacks in this presentation.

In 2001, a new class of decoding algorithms, known as dual attacks, was introduced and has been significantly improved over the past four years. These attacks rely on a reduction of the decoding problem to the Sparse Secret Learning Parity with Noise (sparse-LPN) problem, where samples are constructed using low-weight codewords of the dual (orthogonal) code. As a consequence, dual attacks require, as a core subroutine, the ability to find low-weight codewords in the dual code.

This observation can be seen as a first step toward addressing a long-standing open question: can the decoding problem be dualized, that is, polynomially reduced to a decoding problem in the dual code? This question is of particular importance in code-based cryptography, since a positive answer would imply that decoding at rate R can be reduced to decoding at rate 1 - R. Such a result would be highly significant, as the complexity of existing primal attacks is not symmetric with respect to the code rate.

In this presentation, after a brief overview of primal and dual attacks, we examine the parameter regimes in which dual attacks currently outperform the best known primal attacks. In particular, we focus on the relevance of these attacks for cryptosystems such as BIKE, McEliece, and HQC, where the target decoding distance grows sub-linearly with the code length.

Code Rigidity in Characteristic 2

Author: Adam Downs[1]

[1] *Virginia Tech*

Contact: adamd24@vt.edu

Two linear codes are *equivalent* if there exists a monomial matrix that transforms one to the other. The problem of finding a monomial transformation from one code to another underlies the Linear Equivalence Signature Scheme (LESS). An *automorphism* of a linear code is a monomial matrix which fixes the code. When a code has a large number of automorphisms, it is easier to solve the linear code equivalence problem; thus, it is desirable for cryptography to have codes with as few automorphisms as possible. A code with the smallest possible number of automorphisms is called *rigid*.

Prior results show that almost all binary codes of dimension k and length 2k are rigid as k goes to infinity. We extend this result to arbitrary finite fields of characteristic 2.

Half is Enough: halving keys through optimal representation of self-orthogonal codes

Authors: Marco Baldi[1], Rahmi El Mechri[2], Paolo Santini[1], Riccardo Schiavoni[1]

[1] *Università Politecnica delle Marche*

[2] *Univeristà Politecnica delle Marche, Scuola IMT Alti Studi Lucca*

Contact: rahmi.elmechri@imtlucca.it

Given two linear codes, the Permutation Equivalence Problem (PEP) asks to find a permutation that maps one code onto the other. The state-of-the-art solvers for PEP take time that is either exponential in the code length or in the dimension of the hull, which is the intersection between a code and its dual. To avoid the latter type of attacks, PEP-based cryptosystems employ linear codes with large hull such as self-orthogonal codes, i.e., codes contained in their dual. Due to the need of representing such codes, whose communication cost grows quadratically with their length, such cryptosystems suffer from large public keys. We present an efficient compression technique that allows to represent self-orthogonal codes with fewer bits, taking advantage of the pairwise orthogonality of their codewords. Among the cryptosystems that benefit from this work there are SPECK, LESS (and its threshold variant LEAST), which is a candidate in the ongoing NIST standardization process, and the updatable public key encryption scheme proposed by Albrecht, Benčina and Lai, which is based on a special instance of the Lattice Isomorphism Problem (LIP). Remarkably, the compression technique nearly halves the public key for the former two schemes.

Sample Complexity of the Matrix Code Equivalence Problem

Authors: Wendi Gao, Pedro Juan Soto
Contact: wendig@vt.edu

The Matrix Equivalence Digital Signature (MEDS) is a code-based digital signature that was submitted to the NIST call for quantum-resistant protocols. It is currently considered as a candidate

for building advanced group action signatures schemes.

The hard problem behind this digital signature is the Matrix Code Equivalence problem. Namely, given two matrix codes C_1 and C_2, suppose that there is a matrix pair (A, B) where $AC_1B = C_2$ (left and right multiplication of matrices in C_1 by A and B), how do we find the matrices A and B?

This talk studies the matrix code equivalence problem from a sample complexity point of view. Since there is no known polynomial-time algorithm for finding A and B given one instance of the (C_1, C_2) pair, how many such instances are enough to make a polynomial-time solution possible? In particular, we try to establish a relationship between the number of samples and the complexity of finding the secret pair (A, B) using the XL (Extended Linearization) Algorithm.

Post-Quantum Blind Signatures from Matrix Code Equivalence

Authors: Veronika Kuchta[1], Edoardo Persichetti[1], Jason LeGrow[2]

[1] *Florida Atlantic University*

[2] *Virginia Polytechnic Institute and State University*

Contact: vkuchta@fau.edu

We construct a novel code-based blind signature scheme, using the Matrix Equivalence Digital Signature (MEDS) group action. The scheme is built using similar ideas to the Schnorr blind signature scheme and CSI-Otter, but uses additional public key and commitment information to overcome the difficulties that the MEDS group action faces: lack of module structure (present in Schnorr), lack of a quadratic twist (present in CSI-Otter), and non-commutativity of the acting group. We address security concerns related to public key validation, and prove the security of our protocol in the random oracle model, using the security framework of Kastner, Loss, and Xu, under a variant of the Inverse Matrix Code Equivalence problem and a mild heuristic assumption. We also discuss alternative techniques for constructing a code-based blind signature and consider possible parameter sets and corresponding performance metrics.

Code distances: a new family of invariants of linear codes

Author: Hiram López[1]

[1] *Virginia Tech*

Contact: hhlopez@vt.edu

We introduce code distances, a new family of invariants for linear codes. We establish some properties and prove bounds on the code distances, and show that they are not invariants of the matroid (for a linear block code) or q-polymatroid (for a rank-metric code) associated to the code. By

means of examples, we show that the code distances allow us to distinguish some inequivalent MDS or MRD codes with the same parameters. We also show that no duality holds, i.e., the sequence of code distances of a code does not determine the sequence of code distances of its dual.

An Elementary Analysis of Multivariate Goppa Codes

Author: William Mahaney[1]

[1] *Virginia Tech*

Contact: mahaney1@vt.edu

Goppa codes form a structured family of linear error-correcting codes introduced by Valery D. Goppa in 1970 and later interpreted within the framework of algebraic geometry as codes arising from algebraic curves over finite fields. Binary Goppa codes with irreducible Goppa polynomials are used in the Classic McEliece post-quantum key encapsulation mechanism (PQ-KEM), where their efficient decoding algorithms and resistance to structural attacks underpin a long-standing and well-studied code-based public-key encryption scheme.

Motivated by classical Goppa codes, multivariate Goppa codes are an analogous family of codes constructed using multivariate polynomials, originally introduced by Hiram H. López and Gretchen L. Matthews in 2021. In this talk, we generalize the original construction and show that these generalized multivariate Goppa codes exhibit similar rank and distance bounds, as well as comparable structural properties. We conclude by discussing their connection to classical Goppa codes and their suitability for cryptographic applications.

Relative generalized Hamming weights and their applications in cryptograph

Authors: Rodrigo San-José[1], Diego Ruano[2], Gonzalo Rodríguez-Pajares[2]

[1] *Virginia Tech*

[2] *University of Valladolid*

Contact: rsanjose@vt.edu

The relative generalized Hamming weights of a nested pair of linear codes are a generalization of the minimum distance. We will see how these parameters characterize the security of ramp secret sharing schemes, and how this can be adapted for private information retrieval. The computation of these parameters for a linear code is NP-hard in general, and we will study the most efficient current algorithm to compute them for any nested pair of linear codes.

Some of these results are joint work with Diego Ruano and Gonzalo Rodríguez-Pajares.

Code-Based Arithmetic Circuits

Authors: Freeman Slaughter[1], Felice Manganiello[2]

[1] *University of South Florida*

[2] *Clemson*

Contact: fslaughter@usf.edu

Arithmetic circuits provide a versatile framework for demonstrating generic algebraic statements, expressible as a system of polynomials, in a zero-knowledge manner. While this primitive can be used to prove knowledge of solutions to NP-complete problems (graph 3-coloring, Sudoku, etc), existing implementations generally rely on discrete logarithm problem assumptions. In this talk, we introduce a novel code-based arithmetic circuit framework. Our construction permits a prover to demonstrate that committed Hamming-ball vectors satisfy certain arithmetic relationships, solely by acting on their syndromes. By translating generic circuit satisfiability to this code-based setting, our framework provides a critical stepping stone for the development of a secure code-based cryptocurrency.

MS 15: Linear Algebra Foundations for Data-driven Modeling and Model Order Reduction

Minisymposium Organizers: Ionut Farcas, Steffen Werner

Mathematical models of real-world phenomena are fundamental across numerous scientific and industrial applications, including vibrational analysis, control of mechanical systems, shape and design optimization, control, and the development of digital twins. To be practically useful, these models must achieve high accuracy in representing quantities of interest while remaining computationally efficient. When first-principles models are available and accurately capture the system dynamics, model order reduction techniques offer effective ways to construct low-dimensional, computationally inexpensive surrogate models. Conversely, when the underlying processes are not fully understood, data-driven modeling and system identification methods enable the creation of suitable low-dimensional models directly from data. In both scenarios, concepts from linear algebra are the foundation for the developed computational procedures ranging from the construction of low-dimensional subspaces to the solution of matrix equations to the fitting of nonlinear functions, and many more. This minisymposium will bring together researchers in linear algebra, model order reduction, and data-driven modeling to exchange ideas and discuss the latest advances and challenges in designing low-dimensional surrogate models.

List of Talks

Matrix Structures for Certified Stability of Nonlinear Reduced Models

Author: Peter Benner
Contact: benner@mpi-magdeburg.mpg.de

Learning compact surrogate models from data has become a major application area of machine learning techniques. Such models are required to describe dynamical behavior of processes in the presence of time series data and the absence of explicit mechanistic models. This may be the case if only measurement data is available or simulation data is obtained via proprietary software. Prominent techniques in this area are dictionary learning via sparse regression or operator inference which allow to infer nonlinear models describing the data in some optimal way. If stable behavior of the underlying process is expected, be it Lyapunov or global stability, or some sort of attractor dynamics, the classical system identification approaches can not guarantee this intrinsic property of the physical process. We discuss how several stability concepts can be encoded in the inferred model structure using structured matrices so that the obtained surrogate models are guaranteed to have the desired stability property. We discuss in particular quadratic and cubic nonlinearities as it is known that these are sufficiently expressive for many nonlinear dynamical systems. We also discuss the extension of the suggested encodings to parametric and control systems.

This talk is based on prior work with Pawan K. Goyal, Siddarth Mamidisetti, and Igor Pontes Duff.

Learning magnetic field structure from trajectories

Author: David Bindel[1]

[1] *Cornell University*

Contact: bindel@cornell.edu

Stellarators are non-axisymmetric magnetic field configurations used to confine plasmas. Within a stellarator, particles roughly follow magnetic field lines, and the magnetic fields in stellarators can be organized into different regions according to the dynamics of field line flows, with regions of nested flux surfaces potentially interspersed with islands or regions of chaos. While can visualize these structures with Poincare plots, such visualizations are not wholly satisfactory if we wish to numerically optimize such structures. In this talk, we describe an approach to automatically

inferring the structure of magnetic field line flows by building models of the dynamics of field line trajectories. We give connections to dynamical systems theory, numerical extrapolation methods, and the theory of adaptive filtering.

Joint work with Max Ruth (UT Austin).

Reduced order modeling and numerical linear algebra analogs

Authors: Akil Narayan[1], Amy de Castro[1], Filip Belik[1]

[1] *University of Utah*

Contact: amy.de.castro@utah.edu

Constructing low-order approximations to a high-dimensional manifold is a well-studied field as these types of problems arise naturally from the solution of parametric partial differential equations in multi-query or optimization contexts. Full-order approximations, although the most accurate approach to reconstructing a solution manifold, incur too high of an expense in these scenarios. Results for reduced order modeling (ROM) procedures such as proper orthogonal decomposition (POD) and greedy reduced basis methods are often stated in a continuous, functional analysis setting; however, these algorithms are the continuous analog of well-known discrete linear algebra routines for matrix factorizations. In this talk, we compare widely used ROM techniques with their discrete counterparts: POD with SVD, reduced basis methods with column pivoted QR, and empirical interpolation with full pivoted LU. Results from the continuous and discrete settings are juxtaposed to highlights similarities and allow for the interpretation and development of continuous ROM results in light of their linear algebra decomposition analogs.

Non-intrusive reduced-order models for parameterized partial differential equations using kernel methods

Authors: Alejandro Diaz[1], Anthony Gruber[1], John Tencer[1], Patrick Blonigan[1], Shane McQuarrie[2]

[1] *Sandia National Laboratories*

[2] *Brigham Young University*

Contact: andiaz@sandia.gov

This talk presents an interpretable, non-intrusive reduced-order modeling technique for parameterized problems using regularized kernel interpolation. Parameterized reduced-order models (ROMs) enable the rapid approximation of PDE solutions corresponding to a given parameter, thus accelerating uncertainty quantification or inverse problem workflows requiring many PDE solves. Existing non-intrusive parameterized ROM approaches approximate the ROM dynamics by solving a data-driven least-squares regression problem for low-dimensional matrix operators. However,

these approaches typically assume affine parametric dependence, which may not be satisfied by the underlying full-order model (FOM). To overcome this limitation, our approach leverages regularized kernel interpolation, which yields an optimal approximation of the ROM dynamics from a user-defined reproducing kernel Hilbert space and allows for arbitrary parametric dependence. We further show that our kernel-based approach can produce interpretable ROMs whose structure mirrors the parameterized FOM structure by embedding judiciously chosen feature maps into the kernel. The approach is demonstrated in several numerical experiments.

Numerical linear algebra for data driven nonlinear dynamics

Authors: Zlatko Drmac[1], Igor Mezic[2]

[1] *Faculty of Science, University of Zagreb*

[2] *University of California at Sanata Barbara*

Contact: drmac@math.hr

The Dynamic Mode Decomposition (DMD) is a powerful and versatile numerical method for data driven analysis of nonlinear dynamical systems, with a wide spectrum of applications. It can be used for model order reduction, analysis of latent structures in the dynamics, and e.g. for forecasting and control. The theoretical bedrock upon which the more general Extended DMD (EDMD) framework is built is the Koopman composition operator. The EDMD can be described as a data driven Rayleigh-Ritz extraction of spectral information of a Koopman operator associated with the underlying dynamical system. The nonlinear data snapshots are represented using the eigenvectors of the operator resulting in a modal decomposition KMD (Koopman Mode Decomposition). This becomes a model order reduction tool that represents the nonlinear dynamics using selected eigenpairs. It can be used to reveal coherent states and for forecasting. The numerical realization of the Koopman operator framework (i.e. DMD and EDMD) for computational analysis of nonlinear dynamics is an excellent illustration of the importance and power of numerical linear algebra, and an instructive case study of the software development based on the design principles introduced in the state of the art software packages such as e.g. LAPACK. The main ingredients of all variants of the DMD are the SVD and low rank approximations,the QR factorization and the orthogonal projections, structured least squares approximations, and approximations of eigenvalues and eigenvectors from subspaces. The problem of illconditioned eigenvectors is solved using a Koopman-Schur decomposition, based on unitary transformations. The analysis in terms of the eigenvectors is replaced with a modal decomposition in terms of flag of invariant subspaces that correspond to selected eigenvalues. Accuracy is controlled using computable residuals. We will discuss these and other fine details that are built in robust software solutions.

Tackling the curse of dimensionality through the parametric Loewner framework: recent advances and applications

Authors: Athanasios Antoulas[1], Charles Poussot-Vassal[2], Ion Victor Gosea[3], Pierre Vuillemin[2]

[1] *Rice University, Houston, USA*

[2] *Onera*

[3] *Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany*

Contact: gosea@mpi-magdeburg.mpg.de

We extend the Loewner framework to multivariate (static and dynamic) functions with an arbitrary number n of variables [1]. We present the following facts:
(i) That n-variable rational functions (and realization), described in the barycentric basis, can be constructed to interpolate and/or approximate/compress any tensorized n-D data or n-variate function;
(ii) That these n-variable rational functions can be obtained thanks to a sequence of small-scale single-variable interpolation (performed with Loewner matrices), leading to drastically taming the curse of dimensionality (both in memory and computational effort);
(iii) That such sequence results in the variables decoupling, providing a numerically robust solution to the Kolmogorov Superposition Theorem (KST), restricted to rational functions;
(iv) That the Loewner framework bridges "Approximation theory" (both functions and tensors) with "Systems theory", and provides connections with Kolmogorov Arnold Networks (KAN).

A collection of numerical examples and method comparison illustrates the effectiveness and scalability of the proposed method and its ability to tame the curse of dimensionality [2]. Issues and outlook will be specifically addressed. Some research codes are also presented [3].

[1] https://doi.org/10.1137/24M1656657
[2] https://arxiv.org/abs/2506.04791
[3] https://github.com/cpoussot/mLF

This works is carried out in collaboration with A.C. Antoulas, I-V. Gosea, C. Poussot-Vassal and P. Vuillemin.

Randomized Linear Algebra for Symplectic Model Order Reduction of Hamiltonian Systems

Authors: Robin Herkert[1], Bernard Haasdonk[1] Johannes Rettberg[1], Jörg Fehr[1], Patrick Buchfink[1]

[1] *University of Stuttgart*

Contact: robinherkert@freenet.de

Large-scale Hamiltonian dynamics are governed by

$$\dot{x} = J\nabla H(x), \quad J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix},$$

and arise from spatial discretizations of conservative PDEs as well as in molecular and multibody models. In multi-query, control, and real-time settings, projection-based model order reduction (MOR) is essential, but generic reduced spaces may destroy the symplectic geometry, leading to energy gain or loss and poor long-time behavior. Symplectic MOR therefore seeks a symplectic trial basis $V \in \mathbb{R}^{2n\times 2r}$ with $V^T JV = J_{2r}$, so that the reduced system remains Hamiltonian.

A key offline bottleneck of symplectic MOR is symplectic basis generation from snapshot data, commonly via the complex SVD (cSVD) or the SVD-like decomposition, which can dominate the cost when n and/or the number of snapshots are large.

Building on randomized numerical linear algebra, we present basis construction methods that use random sketching to accelerate the matrix factorizations underlying these basis generation techniques, while preserving symplecticity by construction. In particular, the randomized cSVD (rcSVD) and a randomized SVD-like decomposition yield substantial speedups while retaining accuracy close to their deterministic counterparts in numerical tests [1].

Beyond efficiency, we highlight a priori error analysis for rcSVD-based symplectic MOR. The resulting bounds relate the projection error to sketch size (oversampling) and power-iteration depth, thereby quantifying the runtime–accuracy trade-off, motivating practical parameter choices [2].

References:

[1] R. Herkert et. al. Randomized symplectic model order reduction for Hamiltonian systems. In LSSC 2023, pages 99–107. Springer, Cham, 2024.

[2] R. Herkert et. al. Error analysis of randomized symplectic model order reduction for Hamiltonian systems. LAA, 729:67–99, 2026.

Solution of generalized tensor Lyapunov equations arising in optimal control and model reduction

Authors: Nicholas A. Corbin[1], Boris Kramer[1]

[1] *University of California San Diego*

Contact: bmkramer@ucsd.edu

Hamilton-Jacobi-Bellman partial differential equations (HJB PDEs) arise in various settings in optimal control and model order reduction, and their solutions are notoriously difficult to acquire. For linear time-invariant systems, the HJB PDEs of interest typically simplify to matrix algebraic equations, such as the algebraic Riccati equation or the matrix Lyapunov equation, for which many mature, reliable, and scalable solvers exist. For weakly nonlinear systems with polynomial

dynamics, a popular approach for locally approximating solutions to the HJB PDEs is the method of Al'brekht, which amounts to computing a Taylor series approximation. The first term in the Taylor expansion is given by the solution to a matrix Riccati (or Lyapunov) equation, and then the higher-order terms in the expansion are given by solutions to very large linear algebraic equations. Despite being dense, it turns out that these equations exhibit a surprising amount of structure reminiscent of the matrix Lyapunov equation, and they can be considered as *tensor* Lyapunov equations. In this talk, we expand on related works by considering the generalized tensor Lyapunov equation that arises due to the presence of an invertible but non-identity mass matrix. We describe the changes to the problem associated with the inclusion of the mass matrix, some surprising beauty that emerges in the form of the fractal structure of the Sierpinski triangle, and the resulting obstacles that are introduced that preclude the use of existing solvers, before concluding with a new proposed algorithm for solving these structured systems efficiently. We also highlight the open problems and opportunities for future research on the topic.

Tensor-based Dynamic Mode Decomposition for Complex Dynamics

Author: Feliks Nueske[1]

[1] *Max-Planck-Institute DCTS Magdeburg*

Contact: nueske@mpi-magdeburg.mpg.de

In this talk, I will present recent work on analysing time series data for complex dynamics. Extended dynamic mode decomposition (EDMD, Williams et al, 2015) is a widely used algorithm to learn a linear surrogate model for the statistics of an evolving dynamics, based on the Koopman operator framework. For high-dimensional systems, choosing a suitable basis set can become challenging, as traditional methods are subject to the curse of dimensionality. Tensor-based versions of EDMD were first presented by Klus and co-workers in 2018, and then extended by Nüske et al in 2021. A further extension to approximate the infinitesimal generator of the dynamics was developed by Lücke et al in 2022.

In this talk, I will provide an overview of these developments, and then present some recent results on the application of tensor-based EDMD to simulation data of complex bio-molecules. In particular, I will show how combining dimensionality reduction methods with tensor-based models helps unravel complex structure in high-dimensional time series data.

Approximate Deconvolution and Spatial Filtering of Reduced Order Models for Fluid Flow

Authors: Jorge Reyes[1], Ping-Hsuan Tsai[1]

[1] *Virginia Tech*

Contact: reyesj@vt.edu

This talk focuses on the numerical analysis of regularized projection-based reduced-order models (ROMs) for turbulent fluid flows. Direct numerical simulations are well known to be computationally infeasible for routine simulations in computational fluid dynamics, particularly at high Reynolds numbers. Reduced-order models offer an efficient low-dimensional framework capable of producing fast and reasonably accurate approximations of the full-order dynamics.

For under-resolved flows, which are typical in high-Reynolds-number regimes, standard POD-ROMs often suffer from loss of accuracy and numerical instability. Regularization techniques based on spatial filtering have been shown to improve both stability and accuracy with negligible additional computational cost. However, depending on the choice of filter and regularization strategy, these methods may introduce excessive numerical diffusion, leading to over-smoothing of the resolved dynamics.

To address this issue, we propose an approximate deconvolution approach designed to recover attenuated flow features and mitigate the over-smoothing effects induced by spatial filtering. The resulting framework enhances ROM accuracy while retaining the stabilizing benefits of regularization. Analytical properties and numerical results are presented to demonstrate the effectiveness of the proposed method for turbulent flow simulations.

Reduced Order Modeling of Conservation Laws via CDT

Authors: Aryan Saxena[1], Gustavo Rohde[2], Harbir Antil[1], Ivan V. Medri[2], Kristofor E. Pas[2], Rocio Diaz Martin[3], Sarswati Shah

[1] *George Mason University*

[2] *University of Virginia*

[3] *Florida State University*

Contact: sshah66@gmu.edu

In this work, we propose a reduced-order modeling (ROM) framework for conservation laws that operates in the Cumulative Distribution Transform (CDT) domain. The CDT maps nonnegative, unit-mass states to an L^2 Hilbert space, in which pure translations become affine lines and W_2 (optimal transport) distances coincide with Euclidean distances. This linearization dramatically improves the approximability of advective dynamics: for linear transport, the solution manifold is exactly two-dimensional in CDT space ($d_2 = 0$), and for advection–diffusion we show that, once an initial layer is excluded, the CDT trajectory has uniformly bounded curvature, implying quadratic Kolmogorov n-width decay with n.

Building on these insights, we develop a CDT–POD algorithm: (i) map snapshots to CDT space, (ii) compute a Proper Orthogonal Decomposition (POD) basis in that space, (iii) evolve reduced coefficients via a transform-space surrogate, and (iv) invert the CDT to recover physical states. Numerical experiments on linear advection, advection–diffusion, and both viscous and inviscid

Burgers equations show that CDT–POD attains rank-2 exactness for transported signals, robustly resolves moving shocks and contact discontinuities with orders of magnitude fewer modes than Eulerian POD, and remains competitive in smoothly diffusive regimes. The method preserves mass by construction and is compatible with positivity.

Wavelet-Based Observables for Koopman Analysis: An Extended Dynamic Mode Decomposition Framework

Authors: Cankat Tilki[1], Serkan Gugercin[1]

[1] *Virginia Tech*

Contact: cankat@vt.edu

We present an in-depth analysis of the Koopman operator using wavelet transform. Based on this analysis, we construct an invariant subspace for the Koopman operator and introduce wavelet-based observables that span this subspace. Moreover, on this subspace, we study the Koopman operator and its eigendecomposition. To approximate its action numerically over this subspace, we combine Extended Dynamic Mode Decomposition (EDMD) with the proposed wavelet-based observables, leading to the Wavelet Dynamic Mode Decomposition via Continuous Wavelet Transform (cWDMD) algorithm. We validate our theoretical results on two numerical examples.

Loewner linearizations of structured rational matrices

Authors: Paul Van Dooren[1], Froilan Dopico, Maria del Carmen Quintana, Vanni Noferini

[1] *Universite catholique de Louvain*

Contact: vandooren.p@gmail.com

We show how to use tangential interpolation techniques to construct structured linearizations for several types of structured rational matrices. The classes studied in this paper are square rational matrices that are either Hermitian, or skew-Hermitian, or complex symmetric, or complex skew-symmetric, upon evaluation on one of the following three curves˜: the real axis, the imaginary axis and the unit circle. The proposed linearizations are system matrices for these rational matrices and they preserve the structure of the rational matrices, except for the case of the unit circle. For that case, the rational matrix $R(z)$ is linearized using a palindromic or anti-palindromic system matrix for a modified rational matrix, whose eigenvalues that are not on the unit circle preserve the symmetries of the zeros and poles of $R(z)$. The basic tool used to obtain the results in this paper is tangential interpolation via the Loewner and shifted Loewner matrices. In the case of preserving symmetries with respect to the unit circle, we combine this with Mobius transforms.

Adaptive kernel methods

Authors: Tamas Dozsa, Andrea Angino, Zoltan Szabo, Jozsef Bokor, Matthias Voigt[1]

[1] *UniDistance Suisse*

Contact: matthias.voigt@fernuni.ch

Kernel methods approximate nonlinear maps in a data-driven way by projecting the target map onto a finite-dimensional Hilbert space called the solution space. Traditionally, this space is a subspace of a fixed ambient reproducing kernel Hilbert space (RKHS), determined solely by the chosen kernel and the dataset, whose elements identify the basis elements. Consequently, the projection operator underlying the kernel method depends on the loss function, the dataset, and the choice of ambient RKHS. In this talk, we consider kernel methods whose solution spaces also depend on learnable parameters that are independent of the dataset. The resulting methods can be viewed as variable projection operators that depend on the loss function, the dataset, and the new learnable parameters instead of a fixed RKHS. We propose an efficient approximation of kernels associated with infinite-dimensional RKHSs, commonly used to reduce the solution-space dimension for large datasets. Second, we construct fixed-dimensional, parameter-dependent solution spaces that enable highly efficient kernel models suitable for large-scale problems without the need to approximate kernels of infinite-dimensional RKHSs. Our novel family of adaptive kernel methods generalizes earlier approaches, including Random Fourier Features, and we demonstrate their effectiveness through several numerical experiments.

Dynamical Tensor Train Approximation for Kinetic Equations

Authors: Geshuo Wang[1], Jingwei Hu[1]

[1] *University of Washington*

Contact: geshuo@uw.edu

The numerical solution of kinetic equations is challenging due to the high dimensionality of the underlying phase space. In this paper, we develop a dynamical low-rank method based on the projector-splitting integrator in tensor-train (TT) format. The key idea is to discretize the three-dimensional velocity variable using tensor trains while treating the spatial variable as a parameter, thereby exploiting the low-rank structure of the distribution function in velocity space. In contrast to the standard step-and-truncate approach, this method updates the tensor cores through a sweeping procedure, allowing the use of relatively small TT-ranks and leading to substantial reductions in memory usage and computational cost. We demonstrate the effectiveness of the proposed approach on several representative kinetic equations.

From Structured Loewner Matrices to Balanced Mechanical Systems

Authors: Sean Reiter[1], Steffen W. R. Werner[2]

[1] *Courant Institute of Mathematical Sciences, New York University*

[2] *Virginia Tech*

Contact: steffen.werner@vt.edu

Learning dynamical systems from data has emerged as a pivotal area of research, bridging the realms of mathematics, engineering, and data science. Of particular importance is the construction of models with meaningful internal structure that allows interpretability and explainability of the results. In the case of mechanical and electro-mechanical processes, dynamical systems are typically described via ordinary differential equations with second-order time derivatives

$$M\ddot{x}(t) + D\dot{x}(t) + Kx(t) = Bu(t), \quad y(t) = C_{\mathrm{p}}x(t) + C_{\mathrm{v}}\dot{x}(t),$$

with the system matrices $M, D, K \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, $C_{\mathrm{p}}, C_{\mathrm{v}} \in \mathbb{R}^{p\times n}$, the inputs $u(t) \in \mathbb{R}^m$, the internal states $x(t) \in \mathbb{R}^n$, and the outputs $y(t) \in \mathbb{R}^p$. A popular approach in model order reduction to construct accurate low-dimensional surrogate models that retain the second-order internal structure is the second-order position-velocity balanced truncation method. Thereby, the system is balanced with respect to structured system Gramians and states corresponding to small singular values of these Gramians are truncated.

In this work, we present a new data-driven formulation of the second-order balanced truncation method. Utilizing suitable structured Loewner matrices, we can directly compute low-dimensional balanced second-order systems with generalized proportional damping, $D(s) = f(s)M + g(s)K$ where $f(s), g(s) \in \mathbb{C}$, from given frequency domain measurements. While based on classical structure-preserving model reduction, our method can be used to enforce mechanical model structure independent of the data source. Several numerical examples demonstrate the effectiveness of the proposed method.

MS 16: Hierarchical Low-rank Approximations: Algorithms and Applications

Minisymposium Organizers: Chao Chen, Arvind Saibaba

Large dense matrices arise in a wide range of problems in scientific computing and machine learning. While many such matrices are not globally low-rank, they often contain subblocks that are numerically low-rank. Hierarchical matrix representations exploit this structure to achieve linear or nearly linear complexity in both storage and key computations such as matrix-vector products and the solution of linear systems. This mini-symposium will highlight recent developments in hierarchical low-rank approximations, emphasizing both theoretical advances, practical algorithms and applications. Topics include randomized methods, low-rank approximation, tensor compression, and applications in scientific computing, data analysis, and machine learning.

List of Talks

Estimating High-Dimensional Covariance Matrices with Hierarchical Rank Structure

Authors: Robin Armstrong[1], Samuel Otto[1], Anil Damle[1]

[1] *Cornell University*

Contact: rja243@cornell.edu

Many algorithms in data assimilation and model order reduction rely on sample-based estimates for a covariance matrix associated with the trajectory of a high-dimensional dynamical system. The number of available samples is often far less than the dimension of the underlying state space because of computational constraints. Under these circumstances, extracting meaningful covariance information requires that the noisy statistics of the sample are regularized with a structural assumption. This talk will describe a regularization technique that exploits the rank structure of submatrices representing cross-covariances between well-separated domains of space. We will establish that low-rank truncations of these submatrices can be estimated from fewer samples than the submatrices themselves. We will then show how this fact can be used to encode physics-informed regularizing assumptions onto the sample statistics, resulting in a hierarchically rank-structured covariance estimator. Through numerical experiments with a variety of dynamical systems, we will demonstrate that these techniques are effective at reducing sampling errors in the covariance.

Superfast and stable divide-and-conquer singular value decomposition for hierarchical rank-structured matrices

Authors: Chenyang Cao[1], Jianlin Xia, Xiaofeng Ou

[1] *Purdue University*

Contact: cao302@purdue.edu

This talk gives a superfast divide-and-conquer algorithm for computing the full singular value decomposition (SVD) of hierarchical rank-structured matrices with small off-diagonal ranks. The method achieves nearly linear complexity while delivering all singular values and singular vectors in structured forms. The structured representation of singular vectors enables near-linear operations with reduced storage. In contrast, classical algorithms for the full SVD require cubic computational cost and quadratic storage. The proposed method directly handles nonsymmetric or rectangular matrices by reducing them to a hierarchical block broken arrow form via stable QL factorizations. This form is then repeatedly decomposed through rank-1 SVD updates in the conquering stage. Several stability-preserving mechanisms are incorporated, including deflation, splitting and shifting, and orthogonality-preserving perturbations, to ensure the robustness of this stage. Efficiency is enhanced via fast kernel methods, such as the fast multipole method (FMM). A rigorous numerical error analysis establishes the backward stability of the overall process. Numerical experiments demonstrate significant advantages in both accuracy and efficiency, yielding the first fast and stable SVD for general hierarchical rank-structured matrices.

An adaptive method for constructing hierarchical approximations and its application to inverse problems

Author: Chao Chen[1]

[1] *NC State*

Contact: cchen49@ncsu.edu

Dense matrices arise in many areas of computational science, and hierarchical approximation methods have been developed to reduce their storage and computational costs. However, many existing approaches require access to individual matrix entries, which may not be available or prohibitively expensive to compute in important applications. A prominent example is the Hessian in Bayesian inverse problems: accurate Hessian approximations can enable fast posterior sampling, yet forming Hessian entries is costly, while Hessian–vector products can often be computed efficiently.

In this talk, we present adaptive randomized algorithms for constructing hierarchical approximations to a prescribed accuracy in a matrix-free setting. In particular, we extend the peeling paradigm to an adaptive compression strategy for HODLR matrices by approximating off-diagonal blocks at each hierarchical level to meet a global tolerance. Numerical experiments show that the proposed method achieves high accuracy more efficiently than fixed-rank schemes and than global low-rank approximations of the Hessian.

Parametric Hierarchical Matrix Approximations to Kernel Matrices

Authors: Abraham Khan[1], Arvind Krishna Saibaba[1], Chao Chen[2], Vishwas Rao

[1] *North Carolina State University*

[2] *NC State*

Contact: awkhan3@ncsu.edu

Kernel matrices arising in applications such as Gaussian processes may not always admit a low-rank approximation. Important examples are kernel matrices induced by certain members of the Matérn family of covariance kernels, with smaller length scales and values of ν. Still, they can often be approximated by a hierarchical matrix ($\mathcal{H}$-matrix or $\mathcal{H}^2$-matrix), which consists of a hierarchy of small near-field blocks (sub-matrices) stored in a dense format and large low-rank far-field blocks that are efficiently stored in factored form. A hierarchical matrix approximation of a kernel matrix can be constructed, stored, and used to perform matrix-vector multiplication in log-linear or linear complexity with respect to n. Standard methods for approximating kernel matrices with hierarchical matrices do not account for the following: kernel matrices often depend on certain hyperparameters that must be optimized over a fixed parameter space. For example, in Gaussian processes and Bayesian inverse problems, estimating the hyperparameters from the data involves solving an optimization problem, which requires repeatedly forming or approximating the kernel matrices for a range of parameters. To address this computational challenge, we introduce a new class of hierarchical matrices, namely, parametric (parameter-dependent) hierarchical matrices. The

construction of a parametric hierarchical matrix follows an offline-online paradigm. In the offline stage, the near-field and far-field blocks are approximated by using polynomial approximation and tensor compression. In the fast online stage, for a particular hyperparameter, the parametric hierarchical matrix is instantiated efficiently as a standard hierarchical matrix. Numerical experiments show speedups of over 100 times compared with existing techniques.

Extension of hierarchical matrices to hierarchical tensors with butterfly compression

Author: Yang Liu[1]

[1] *Lawrence Berkeley National Laboratory*

Contact: liuyangzhuan@lbl.gov

The development of hierarchical matrix techniques has been essential for many modern scientific computing frameworks including fast direct solvers for PDEs and integral equations, scalable kernel methods and Gaussian processes, and second-order optimization and inverse problems, etc. These algorithms oftentimes lead to optimal computational and memory complexities. That said, when dealing with higher-dimensional problems, e.g., PDEs and kernels in at least 3D spaces, matrix algorithms can still be expensive and tensor algorithms become more attractive. Despite several recent pioneering works attempting to extend hierarchical matrices to hierarchical tensors, the field of hierarchical tensors remains largely unexploited. This work attempts to tackle the computational challenge of fast tensor algorithms for highly-oscillatory and high-dimensional integral operators. We develop hierarchical tensors with tensor butterfly compression, and demonstrate its efficiency for computing Green's function tensors in volume integral equations and Babich ansatz for high-frequency wave equations. This is a joint work with Jianliang Qian, Yuxiao Wei, Tianyi Shi, Hengrui Luo and Paul Michael Kielstra.

Circulant Preconditioning fractional PDEs on Adaptive Meshes

Authors: Kate Wall[1], James Adler, Misha Kilmer[1], Xiaozhe Hu[1]

[1] *Tufts University*

Contact: kate.wall@tufts.edu

A preconditioner for solving fractional partial differential equations (PDEs) is presented. In our method the fractional PDE is discretized on an adaptive grid, resulting in a Hierarchical matrix representation. The stiffness matrix has Toeplitz blocks along the diagonal and low-rank approximations off the diagonal. Our preconditioner expands on previously developed methods of conditioning Toeplitz systems with circulant matrices. We show how these methods can be applied cheaply on the adaptive mesh and prove that the spectrum of the resulting system is well-clustered. In order

to prove these results, we must take special consideration of how the low-rank blocks perturb the eigenvalues of the Toeplitz block-diagonal system. We validate our results for various fractional orders and inspect any assumptions through numerical tests.

Unknown hierarchies, hyperbolic PDE, and randomized rank detection

Authors: Christopher Wang[1], Alex Townsend[1]

[1] *Cornell University*

Contact: cyw33@cornell.edu

We describe a problem arising from operator learning for hyperbolic PDEs, in which one would like to recover an unknown, non-standard low-rank hierarchical partition of a linear operator using only input-output data, or, in the finite-dimensional case, matrix-vector products. We provide a solution to the operator learning problem by employing a continuous analogue of the randomized SVD (RSVD) to decide whether the operator, restricted to a given subdomain, is numerically low-rank or not. Doing so requires the RSVD to obtain good singular subspace estimates, which in theory depends on the sizes of gaps between singular values of the operator. Thus, in the second part of this talk, we derive exact descriptions for the angular error of the approximate singular subspaces returned by the RSVD, which helps explain why large singular value gaps are typically not required in practice.

Adaptive coarse space for multi-level overlapping Schwarz preconditioners in FROSch

Authors: Jascha Knepper[1], Alexander Heinlein[2], Ichitaro Yamazaki[3]

[1] *University of Cologne*

[2] *TU Delft*

[3] *Sandia Labs*

Contact: iyamaza@sandia.gov

We discuss the adaptive coarse-space basis functions for the multi-level overlapping additive Schwarz preconditioners implemented in FROSch. The basis functions are formed based on the discrete Harmonic extensions of the local subdomain interface functions. The basis functions for the interface are composed of the eigenvectors corresponding to the small eigenvalues of the generalized eigenvalue problems formed by the local Schur complements on the interface. This leads to a provably-robust multi-level preconditioner for solving heterogeneous elliptic problems. We present numerical results to demonstrate that the iterations count stays relatively constant even with the increasing number of subdomains for solving problems with large coefficients jumps. Our performance results on

NERSC Perlmutter supercomputer will then demonstrate the scalability of the method for solving large-scale problems.

MS 17: Where Algebraic Coding Theory and Graph Theory Meet

Minisymposium Organizers: Aida Abiad, Giuseppe Cotardo

This session highlights the deep and evolving connections between algebraic coding theory and graph theory, a relationship that traces back to Delsarte's foundational work and has grown through decades of shared linear-algebraic and combinatorial ideas. We focus on how graph-theoretic techniques, algebraic invariants, and combinatorial constructions inform modern code design, while coding-theoretic perspectives continue to influence developments in algebraic and network graph theory. Topics include extremal and spectral methods for algebraic and combinatorial constructions, graph-based interpretations of code metrics, and new applications in network communication, distributed computing, and secure information sharing.

List of Talks

A graph from the injection metric

Authors: Aida Abiad[1], Alberto Ravagnani[1], Matteo Bertuzzo[1]

[1] *Eindhoven University of Technology*

Contact: m.bertuzzo@tue.nl

The injection metric measures the distance between two subspaces and naturally arises in the context of subspace coding, when the codewords potentially have different dimensions. In this talk,

we consider a graph associated with the injection metric and examine how its structure determines the existence of good codes for subspace coding.

Secure Distributed Matrix Multiplication

Author: Rafael D'Oliveira[1]

[1] *Clemson University*

Contact: rdolive@clemson.edu

Matrix multiplication is, oftentimes, the most expensive computational task in an algorithm. It is the computational bottleneck for training many of the now well-celebrated learning algorithms, for example. To speed up the algorithm, the data can be distributed on many machines to perform the computations in parallel. This sharing of the data, however, raises security concerns when the data is sensitive and has to remain private, such as financial or medical data. Secure distributed matrix multiplication (SDMM) studies how to parallelize matrix multiplication while keeping the data secure.

Graph-based error correction code constructions made practical by modern decoder developments

Authors: Ken Duffy[1], Lukas Rapp[2], Muriel Medard[2]

[1] *Northeastern University*

[2] *MIT*

Contact: k.duffy@northeastern.edu

The study of error correcting codes has two important facets: code construction and decoder design. Many graph-based code constructions have been established to have desirable theoretical properties, but, heretofore, have not been practically decodable. In this talk, we explain recent developments in code-agnostic decoders, including soft-input soft-output variants, that offer a way forward in the provision of a practical decoder for such codes.

Hierarchical quasi-cyclic codes from polynomial evaluation codes

Authors: Emily McMillon[1], Kathryn Haymaker[2]

[1] *Rice University*

[2] *Villanova University*

Contact: kathryn.haymaker@villanova.edu

In this talk, we present a family of algebraically constructed hierarchical quasi-cyclic codes. These codes are built from Reed-Solomon and polynomial evaluation codes using a construction of superimposed codes by Kautz and Singleton. Using a novel ordering of the codewords and evaluation points, we show both the number of levels in the hierarchy and the index of these q-ary-derived codes are determined by the field size. We compute explicit code parameters and properties as well as some additional bounds on parameters such as rank and distance. In particular, starting with Reed-Solomon codes of dimension two yields hierarchical quasi-cyclic codes with Tanner graphs of girth 6. Finally, we present a table of small code parameters and note that some of these codes meet the best known minimum distance for binary codes, with the additional hierarchical quasi-cyclic structure.

On the Hamming Weight Distribution of Cyclic Codes with Arbitrarily Many Nonzeroes

Author: Shuxing Li[1]

[1] *University of Delaware*

Contact: shuxingl@udel.edu

Determining the minimum Hamming distance of an error-correcting code $\mathcal{C}$ has long stood as a fundamental challenge in coding theory. In this talk, we turn our attention to an even more ambitious problem: determining the full Hamming weight distribution of $\mathcal{C}$. This means counting the number of codewords in $\mathcal{C}$ of each possible Hamming weight—an essential but notoriously difficult task, especially for general codes.

While computing the weight distribution for an arbitrary code remains largely out of reach, the problem becomes more tractable under additional structural assumptions, especially when the code $\mathcal{C}$ is cyclic. We will discuss the calculation of three families of cyclic codes with arbitrarily many nonzeroes, highlighting its connections to linear algebra, graph theory, and association schemes.

The talk is based on joint work with Gennian Ge (Capital Normal University), Sihuang Hu (Shandong University), Tao Feng (Zhejiang University); as well as Maosheng Xiong (Hong Kong University of Science and Technology) and Haode Yan (Harbin Institute of Technology, Shenzhen).

An Eventown Result for Permutations

Author: Nathan Lindzey[1]

[1] *University of Memphis*

Contact: nathan.lindzey@memphis.edu

A family of permutations $\mathcal{F} \subseteq S_n$ is *even-cycle-intersecting* if $\sigma\pi^{-1}$ has an even cycle for all $\sigma, \pi \in \mathcal{F}$. We show that if $\mathcal{F} \subseteq S_n$ is an even-cycle-intersecting family of permutations, then $|\mathcal{F}| \leq 2^{n-1}$, and that equality holds when n is a power of 2 and $\mathcal{F}$ is a double-translate of a Sylow 2-subgroup of S_n. This result can be seen as an analogue of the classical eventown problem for subsets and it confirms a conjecture of János Körner on maximum reversing families of the symmetric group. Along the way, we show that the canonically intersecting families of S_n are also the extremal *odd-cycle-intersecting* families of S_n for all even n. While the latter result has less combinatorial significance, its proof uses an interesting new character-theoretic identity that might be of independent interest in algebraic combinatorics.

External codes for multiple unicast networks via interference alignment

Authors: Alberto Ravagnani[1], Felice Manganiello[2], Frank Kschischang[3], Kristen Joyce

[1] *Eindhoven University of Technology*

[2] *Clemson University*

[3] *Univerisity of Toronto*

Contact: manganm@clemson.edu

We introduce a formal framework to study the multiple unicast problem for a coded network in which the network code is linear over a finite field and fixed. We show that the problem corresponds to an interference alignment problem over a finite field. In this context, we establish an outer bound for the achievable rate region and provide examples of networks where the bound is sharp. We finally give evidence of the crucial role played by the field characteristic in the problem.

Graph Properties of Codes from Dyadic and Quasi-Dyadic Matrices

Authors: Anthony Gómez-Fonseca[1], Gretchen Matthews[2], Kirsten Morris[2], Tefjol Pllaha[1]

[1] *University of South Florida*

[2] *Virginia Tech*

Contact: kdmorris@vt.edu

Dyadic matrices are a subclass of matrices known as reproducible matrices, where the entries of the matrix are completely determined by their first row. Quasi-dyadic matrices are block matrices with dyadic matrices in the blocks.

There has been extensive work analyzing quasi-cyclic codes, codes defined by quasi-cyclic parity check matrices, but less is known about codes arising from dyadic or quasi-dyadic matrices. In this

work, we present results on the properties of the Tanner graphs of dyadic and quasi-dyadic codes. In particular, we present results on their isomorphism classes, absorbing sets, and conditions for constructing quasi-dyadic matrices with Tanner graphs of a certain girth. We also compare the performance of these code constructions under Belief Propagation decoding.

Modern Expander-Based Error-Correcting Codes

Author: Pedro Paredes[1]

[1] *Princeton University*

Contact: pparedes@cs.princeton.edu

This talk explores the interplay between coding theory and expander graphs. We will discuss key developments in the design of expander-based codes, including recent advancements that lead to locally testable codes and efficient quantum codes. This includes a review of work dedicated to designing vertex expanders and unique-neighbor expanders, focusing on the specific properties that make them suitable for modern error correction.

Spectral analysis of linear codes

Authors: Adrián Fidalgo-Diaz[1], Gonzalo Rodríguez Pajares[1], Valentino Smaldore[2]

[1] *Universidad de Valladolid*

[2] *Università degli Studi di Padova*

Contact: valentino.smaldore@unipd.it

Let $\mathcal{C} = \{c_0, c_1, \ldots, c_{q^k}\} \subseteq \mathbb{F}_q^n$ be a $[n,k]_q$-linear code endowed with the Hamming metric. That is, $\mathcal{C}$ is a k-subspace of $\mathbb{F}_q^n$. Let $M_{\mathcal{C}} \in \mathbb{R}^{q^k \times q^k}$ be the distance matrix of the code defined as $(M_{\mathcal{C}})_{i,j} := d_H(c_i, c_j) = |\{i\colon x_i \neq y_i\}|$. We analyze the spectrum of $M_{\mathcal{C}}$, providing examples of spectra for well-known families of linear codes. We compute the multiset of the eigenvalues of $M_{\mathcal{C}}$ and bases for their respective eigenspaces.

Coding theory via graph theory

Author: Michael Tait[1]

[1] *Villanova University*

Contact: michael.tait@villanova.edu

We discuss the use of graph theory methods to attack coding theory problems.

This is joint work with Aida Abiad and Harper Reijnders.

MS 18: Topics in Randomized Numerical Linear Algebra

Minisymposium Organizers: Jamie Haddock, Anna Ma, Kate Pearce

Randomized Numerical Linear Algebra (RandNLA) uses randomization to make algorithms for large-scale linear algebra problems faster, more memory-efficient, and often simpler, while still giving high-accuracy results with high probability. This minisymposium will serve as a follow-up event to the Research Collaboration Workshop, "Randomized Numerical Linear Algebra" (RNLA), held at the Institute of Pure and Applied Mathematics (IPAM) in August 2025. Topics that will be featured in this session include (i) randomized algorithms for solving inverse problems in X-ray science, (ii) structure-aware randomization for linear algebra, (iii) Randomized Krylov Methods for Large-scale Inverse Problems, and (iv) randomization in transformer models, with speakers from the workshop.

List of Talks

Robust hybrid infinite and finite dimensional tensor factorizations

Authors: Jocelyn Chi[1], Eric Chi[1]

[1] *University of Minnesota Twin Cities*

Contact: jocetchi@gmail.com

The canonical polyadic (CP) tensor decomposition represents a multidimensional data array as a sum of rank-one outer products of latent factors. Building on CP-HiFi, the hybrid infinite- and finite-dimensional CP framework of Larsen et al. (2024), which introduces quasitensors by modeling selected modes as smooth functional factors in a reproducing kernel Hilbert space, we replace the

standard least-squares objective with a quantile-based loss. This modification enables estimation of conditional quantile functions (rather than conditional means) along functional modes. The resulting approach yields a distribution-aware tensor factorization that is robust to outliers and captures heterogeneity beyond the mean structure.

Accelerating Asynchronous Iterative Methods with Residual-Biased Randomization

Author: Evan Coleman[1]

[1] *University of Mary Washington*

Contact: ecolema4@umw.edu

Asynchronous iterative methods, such as Asynchronous Jacobi, offer a promising mechanism for overcoming synchronization bottlenecks in massively parallel and heterogeneous computing environments. By allowing processing elements to update components using the latest available data without global barriers, these methods maximize computational throughput. However, asynchrony also makes performance more sensitive to the order in which unknowns (or blocks of unknowns) are updated, since processors may act on delayed information and progress unevenly. While randomized numerical linear algebra has provided bounds for asynchronous solvers using uniform sampling, this approach ignores the local error distribution, potentially wasting updates on converged components.

We consider a family of sampling strategies that range from uniform randomization to non-uniform importance sampling, and then to dynamic, time-varying distributions that adapt online using inexpensive residual-based indicators. The main theoretical focus is on convergence in expectation under randomized update rules. To make the analysis transparent, we begin with a synchronous randomized model and establish conditions under which the associated sampling schemes yield expected decrease in a suitable error measure. We then carry the same randomization patterns to an asynchronous setting, showing how mild assumptions on delays and update coverage allow the expected contraction arguments to be extended, while making explicit where asynchrony changes the requirements and the bounds.

Accelerating Ptychographic Reconstruction via Stochastic Multilevel Optimization

Authors: Zichao Wendy Di[1], Borong Zhang[2], Qin Li[2]

[1] *Argonne National Lab*

[2] *UW-Madison*

Contact: wendydi@anl.gov

Ptychography is a powerful coherent diffraction imaging technique essential for reconstructing high-resolution, complex-valued images from intensity-only measurements. However, the reconstruction poses significant challenges due to its nonconvex and ill-posed nature. We propose a novel multilevel optimization framework emphasizing stochastic learning principles to efficiently address these challenges. By formulating the inverse problem as iterative minimization of quadratic surrogate models at varying resolution levels, our method strategically utilizes stochastic gradient evaluations, significantly reducing computational overhead. The multilevel structure exploits hierarchical, multi-scale information intrinsic to the imaging data, enabling faster and more stable convergence compared to traditional deterministic approaches. Hyperparameters are automatically adjusted across resolution levels, ensuring robustness and scalability. Numerical results demonstrate that our stochastic multilevel optimization substantially enhances both reconstruction accuracy and computational efficiency, presenting a robust solution for large-scale ptychographic imaging applications.

Randomized Numerical Linear Algebra for Tensor-Based Transformers

Authors: Katherine Pearce[1], Anna Ma[2], Elizaveta Rebrova[3]

[1] *Oden Institute, UT Austin*

[2] *UC Irvine*

[3] *Princeton*

Contact: katherine.pearce@austin.utexas.edu

Attention mechanisms are a central component of transformer models that capture contextual relationships between tokens in large language models. Although many of the underlying computations (e.g., query, key, and value embeddings in multi-head attention) are inherently multi-way, classical transformer models are built on matrix-based formulations.

In this talk, we discuss several ways that tensorial structure can be imposed on and exploited in attention mechanisms of transformer models. We describe how tensor-based attention can capture higher-order contextual relationships among tokens, vs. pairwise or dot-product attention. We then explore how randomized algorithms in numerical linear algebra may be used to accelerate tensor-based attention computations and reduce storage requirements.

Blocked Leverage Score Sampling in the Randomized Alternating Least Squares CP Tensor Decomposition

Authors: Jamie Haddock[1], Julian Bellavita[2], Kamrun Nahar Keya[3], Karl Pierce[4], Tamara Kolda[5], Yucong Liu[6]

[1] *Harvey Mudd College*

[2] *Cornell University*

[3] *Arizona State University*

[4] *University of Maryland, College Park*

[5] *MathSci.ai*

[6] *Georgia Institute of Technology*

Contact: karl.m.pierce@gmail.com

The CANDECOMP/PARAFAC (CP) decomposition is a powerful tool used for multiway data analysis and to break the "curse of dimensionality" associated with higher-order tensors. The most common way to compute the CP decomposition of a tensor is via a standard alternating least squares (CP-ALS) algorithm. With the CP-ALS, one must iteratively solve a set of overdetermined least squares problem which can quickly become expensive in time and memory. Recently, Battaglino et al[1] developed a randomized CP-ALS algorithm that significantly reduced the cost of the optimization procedure with little impact on the CP decomposition's accuracy. The strategy leverages the structure of the Khatri-Rao product to efficiently compute a sampling probability distribution for each least squares sub-problem. Unfortunately, on modern computer architecture sampling a tensor is a memory-bound operation, i.e. the computational performance of the randomized CP-ALS method can become bound by the bandwidth and latency of a machine. This work seeks to improve on the computational performance of the randomized, sampled CP-ALS strategy by studying the effects of sampling blocks of columns in each least squares problem. In this talk, we illustrate how block-wise sampling impacts the theoretical performance of the randomized CP-ALS algorithm and provide numerical results comparing the block and standard sampled CP-ALS decomposition.

References:
[1]. Casey Battaglino, Grey Ballard, and Tamara G. Kolda. A practical randomized CP tensor decomposition. SIAM J. Matrix Anal. Appl., 39(2):876–901, 2018.

Block Subset Selection based on Randomized QR with Column Pivoting

Authors: Mitchell Scott[1], Irene Simó Muñoz[2], Jamie Haddock[3], Riley Murray[4], Tanya Tafolla[5]

[1] *Emory University*

[2] *Cornell University*

[3] *Harvey Mudd College*

[4] *Sandia National Laboratory*

[5] *University of California, Merced*

Contact: mitchell.scott@emory.edu

The column subset selection problem seeks to find a collection of the matrix columns that have similar spectral properties to the original matrix. Recently with the large amount of data available, many have turned to using randomization to reduce the problem's computation. While there have been many methods that motivate how to select these columns, they are just that–individual columns. However, by blocking these columns, there is less computational communication needed, which makes the process faster. In this talk we will discuss optimality conditions for selecting these block of columns using randomization. We relate the worst case of this randomized method to the deterministic block QR with column pivoting (QRCP). We then corroborate this analysis with numerical experiments to showcase the method's speed and accuracy compared to standard QRCP algorithms.

Singular values and vectors of sparse random rectangular matrices at criticality

Authors: Haixiao Wang[1], Ioana Dumitriu[2]

[1] *University of Wisconsin-Madison*

[2] *University of California, San Diego*

Contact: hwang2594@wisc.edu

In modern machine learning applications, data matrices are always assumed to admit the *signal-plus-noise* structure. Typically, we assume that the spectra of signal and noise matrices are well-separated and that noise subspaces only produce a marginal influence. While these assumptions are readily verified for dense matrices via classical random matrix theory, real-world data is often sparse, posing significant theoretical challenges. We investigate the spectral properties of a matrix $X \in \mathbb{R}^{n\times m}$ with i.i.d. Bernoulli(p) entries. Previous literature proved that when $np \gg \log(n)$, the singular values of X almost surely remain within the compact support of the Marčenko-Pastur (MP) distribution. However, we identify a critical sparsity regime $p = b\log(n)/\sqrt{mn}$ where this classical result fails. We provide a quantitative characterization of the emergence of outlier singular values. For explicit thresholds b_* and b^* as functions of the aspect ratio $\gamma = n/m \geq 1$, we prove a three-phase transition: (1) for $b > b_*$, no outliers exist; (2) for $b^* < b < b_*$, outliers emerge only beyond the right edge of the MP law; and (3) for $b < b^*$, outliers appear on both sides of the bulk, all with high probability. The locations of those outliers are precisely determined by the largest and smallest degree vertices of the underlying random graph. Besides, behavior of singular vectors corresponding to bulk and edge singular values can be characterized precisely. Our approach follows the framework established by Alt, Ducatez, and Knowles (2021), which can be extended to sparse matrices with general bounded entries.

MS 19: Convex Structures in Quantum Information and Gravity

Minisymposium Organizer: Martin Plávala

Convex structures are pivotal in quantum information theory: the convex structure of the cone of positive semidefinite matrices directly impacts the security of quantum key distribution protocols, computational speedups provided by quantum computers, and also plays a crucial role in recently-proposed experiments to test the quantum properties of gravity. Beyond positive semidefinite operators, other convex structures appear repeatedly in quantum theory: the convex set of local and no-signaling behaviors is crucial for our understanding of Bell experiments, various tensor products constructed as symmetric monoiodal categories of ordered vector spaces are often key in expressing and proving practical results, and the properties of convex sets of positive and completely-positive linear maps are tied to several major conjectures in quantum information. This minisymposium aims to bring together researchers specializing in the mathematical, algebraic, and convex aspects of quantum theory in order to share recent results and discuss future applications of cutting-edge mathematical results in applications of quantum theory, all the while making the topics accessible to other mathematicians. By sharing recent insights and exploring future applications, we aim to foster collaboration and inspire innovative research that harnesses cutting-edge mathematical techniques to advance quantum science.

List of Talks

Quantum Simulation of Non-Positive-Operator-Valued Measurements in General Probabilistic Theories with Post-Selection and Prior Information

Authors: Hayato Arai[1], Masahito Hayashi[2]

[1] *University of Tokyo*

[2] *Chinese University Hong Kong, Shenchen*

Contact: h.arai6626@gmail.com

In quantum theory, POVMs form the maximal class of measurements compatible with the Born rule. Operational reconstructions motivate a broader convex framework—General Probabilistic Theories (GPTs)—specified by a convex state space and its dual cone of affine measurement functionals.

Within GPTs one can define "non-positive POVMs" (N-POVMs): Hermitian effects summing to the unit but not positive semidefinite. Although invalid on the full quantum state space, they are mathematically well-defined on suitable convex domains where all outcome probabilities are nonnegative and can outperform genuine POVMs on selected state families (e.g., in state discrimination).

Are such N-POVMs physically realizable? We give a constructive quantum simulation: for a given N-POVM we build a quantum POVM with a failure outcome whose success-conditioned statistics reproduce the N-POVM on an implementation domain. The simulation has two distinct costs—post-selection (success probability) and domain restriction, interpreted as prior information about admissible states—yet the domain can still contain the families exhibiting an N-POVM advantage, so the separation from POVMs persists under the constraints of the implementation. This provides a quantitative bridge between convex operational theories and quantum mechanics.

Time-evolution in generalized probabilistic theories

Authors: Ismael Septembre[1], Joel Huber[2], Martin Plávala[3], Matthias Kleinmann[2], Stefan Nimmrichter[1]

[1] *University of Siegen*

[2] *University of Münster*

[3] *Leibniz University Hannover*

Contact: matthias.kleinmann@uni-siegen.de

Generalized probabilistic theories (GPTs) are a general framework to describe physical theories like quantum mechanics and classical mechanics. At their core, GPTs model prepare-and-measure scenarios by describing preparations and measurements within ordered vector spaces and then predicting measurement outcome probabilities. This framework has been very successful in describing and understanding correlations predicted by quantum theory. More recently, the attention has also

been extended to the description of Markovian time evolution within GPTs by using positivity-preserving dynamical maps. This extension is important from a fundamental physics perspective because, for example, hypothetical non-quantum dynamics induced by gravity would be detectable via precise tests of the time evolution. Furthermore, dynamical GPTs can also be used in order to better understand quantum dynamics and may enable, for example, quantum computational advantage. In this talk, I will discuss the structure of dynamical GPTs, how they can be used to perform precision tests of quantum physics and help to understand dynamical quantum effects that may serve as precursors to applications in quantum technologies.

Information-theoretic finite de Finetti theorems for quantum and beyond

Author: Gereon Koßmann[1]

[1] *RWTH Aachen University*

Contact: kossmann@physik.rwth-aachen.de

We develop quantitative de Finetti representation theorems beyond standard quantum settings, driven by the principle that permutation symmetry enforces approximate independence at finite extension level. First, using a GPT-motivated notion of relative entropy (via an integral representation) we define mutual information for general convex state spaces and prove a uniform monogamy bound for multipartite extensions: for permutation-invariant states, the total mutual information is bounded by a constant depending only on the A-system. This yields an information-theoretic finite de Finetti theorem for convex bodies, asserting that the AB-marginal of an n-extendible (max-tensor) state is close to a separable (min-tensor) state. Second, for constrained separability problems arising in quantum information, we establish constrained de Finetti theorems compatible with additional linear marginal/fixed-point constraints, including a Bose-symmetric variant that operates directly on symmetric-subspace-supported extensions. These results provide quantitative trace-norm closeness of constrained symmetric (or Bose-symmetric) extensions to convex mixtures of constrained product states, enabling de Finetti control in settings where extremal decompositions into pure products are unavailable. The work is based on arXiv:2507.12326, 2507.12302 and 2601.15184.

Today's Experiments Suffice to Indirectly Verify the Quantum Essence of Gravity

Author: Martin Plávala[1]

[1] *Leibniz University Hannover*

Contact: martin.plavala@itp.uni-hannover.de

The gravity-mediated entanglement experiments employ concepts from quantum information to argue that if entanglement due to gravitational interaction is observed, then gravity cannot be described by a classical system. However, the proposed experiments remain beyond out current technological capability, with optimistic projections placing the experiment outside of short-term future. Here we argue that current matter-wave interferometers are sufficient to indirectly prove that gravitational interaction creates entanglement between two systems. Specifically, we prove that if we experimentally verify the Schrödinger equation for a single delocalized system interacting gravitationally with an external mass, then, under one of two reasonable assumptions, the time evolution of two delocalized systems will lead to gravity-mediated entanglement.

Definitely Not That One: The Art of Antidistinguishing Quantum States

Author: Jamie Sikora[1]

[1] *Virginia Tech*

Contact: sikora@vt.edu

Identifying quantum states is one of the oldest problems in quantum information theory. In this work, we explore a variation of this task: rather than determining which state a system is in, we seek to identify a state that it is not. A set of quantum states is said to be antidistinguishable if this inverse guessing game can be won with certainty. We establish tight bounds characterizing when a set of pure states must be antidistinguishable, as well as when it is impossible.

This is joint work with Vincent Russo, Nathaniel Johnston, and Benjamin Lovitz.

Comparing measurement incompatibility via convex subsets of states

Author: Ryo Takakura[1]

[1] *University of Osaka*

Contact: takakura.ryo.qiqb@osaka-u.ac.jp

In quantum theory, certain observables cannot be measured simultaneously, a feature known as measurement incompatibility. This concept captures a fundamental limitation of quantum measurements and has deep connections to quantum phenomena. In this talk, we propose an operational framework to characterize measurement incompatibility using restricted sets of states, modeled as convex subsets. We introduce an operational quantification and ordering defined by the detectability of incompatibility over these subsets. We show that these notions establish a fine-grained hierarchy for measurement incompatibility, providing new structural insights into mutually

unbiased qubit observables beyond standard noise robustness measures.

Convergent Inner-outer Approximation Schemes From De Finetti Theorems For Games And Quantum Error Correction

Author: Julius Alexander Zeiss[1]

[1] *RWTH Aachen*

Contact: julius.zeiss@icloud.com

We apply information-theoretic de Finetti principles to build convergent approximation schemes with explicit finite-level guarantees, yielding both outer relaxations and certified inner points. For polynomial optimization over convex bodies with local equality and inequality constraints, an information-theoretic monogamy argument yields a convergent conic hierarchy whose approximation error decays as $\mathcal{O}[1/\sqrt{n}]$ with the extension level, resolving the lack of finite-level guarantees for inequality-constrained de Finetti-based methods. A constructive rounding scheme converts outer feasible points into certified interior approximations, producing matching inner/outer bounds that "sandwich" the true optimum. As an application, we express the optimal GPT value of two-player non-local games as a polynomial optimization problem with local constraints, enabling finite-convergence approximation guarantees in the GPT setting. In the quantum setting of fixed-size free games with bounded entanglement assistance, we obtain SDP outer hierarchies with ε-additive guarantees and combine them with measurement-based rounding to generate inner sequences of feasible strategies. Crucially, we achieve overall poly$[1/\varepsilon]$ complexity by exploiting representation-theoretic symmetry reduction (symmetric subspace/Bose symmetry, Schur–Weyl duality, block decompositions) to construct the SDPs directly in symmetry-adapted bases. Finally, for approximate quantum error correction under symmetric noise, we round extendability-based outer bounds into explicit encoder–decoder pairs and develop a framework that combines noise symmetries with extendability symmetry (via commuting group actions) to make higher hierarchy levels computationally feasible. The work is based on arXiv:2507.12326, 2507.12302 and 2601.15184.

MS 20: Eigenvalues of Nonnegative and Stochastic Matrices

Minisymposium Organizer: Brando Vagenende

Stochastic matrices are fundamental objects in matrix theory, probability, and dynamical systems, arising naturally in e.g. Markov chains. Their eigenvalues determine convergence rates, stability, periodicity, and long-term behavior, making spectral analysis a central tool in both theory and applications. This minisymposium focuses on developments in the study of eigenvalues and eigenvalue regions of nonnegative matrices in general, with special focus on stochastic matrices and important subsets such as doubly stochastic matrices and monotone stochastic matrices. Topics can include inverse eigenvalue problems, spectral questions for subsets of nonnegative matrices, properties of the Karpelevich arcs, alternative characterisations of the Karpelevich region, characterisation of admissible spectra for stochastic subsets, eigenvalues and geometry, ... The goal of the minisymposium is to bring together researchers working on theoretical, computational, and applied aspects of stochastic matrix spectra and, more broadly, the spectral theory of nonnegative matrices, to highlight new techniques and results, and to stimulate discussion on emerging open problems in the spectral theory of nonnegative matrices and their subsets.

List of Talks

New Results on the Doubly Stochastic Inverse Eigenvalue Problem

Author: Ludovick Bouthat[1]

[1] *Université Laval*

Contact: lubou309@ulaval.ca

Stochastic matrices are matrices with nonnegative entries whose rows each sum to 1. When a matrix and its transpose are both stochastic, it is said to be \emph{doubly stochastic}. In 1938,

Kolmogorov proposed the problem of characterizing the region of possible eigenvalues of an $n \times n$ stochastic matrix, and Karpelevich gave a complete description thirteen years later. This talk concerns the doubly stochastic analogue: characterizing the region ω_n of eigenvalues of $n \times n$ doubly stochastic matrices, which is contained in the unit disk.

Perfect and Mirsky (1965) conjectured that ω_n is the union of the regions Π_k (the convex hulls of the k-th roots of unity) for $k = 1, \ldots, n$. This conjecture holds for $n = 1, 2, 3, 4$, but fails for $n = 5$. The case $n \geq 6$ remains open. In response to the scarcity of progress over the past 60 years, Levick, Pereira, and Kribs proposed a related conjecture. In this talk, I prove a stronger version of this conjecture.

The approach, based on majorization theory and geometric properties, also provides a potential general framework for characterizing ω_5, as well as a numerical method for testing the conjecture for various values of n.

C-realizability in the Symmetric Nonnegative Inverse Eigenvalue Problem: a combinatorial characterization

Authors: Julio Moro[1], Carlos Marijuán[2]

[1] *Universidad Carlos III de Madrid*

[2] *Universidad de Valladolid*

Contact: jmoro@math.uc3m.es

The Real Nonnegative Inverse Eigenvalue problem (hereforth, RNIEP) consists, for a given positive integer n, in characterizing the lists of n real numbers which are the spectrum of some n × n matrix with real entries. C-realizability was originally introduced in [1] as a sufficient condition for the RNIEP. It was shown back then that C-realizability included as particular cases most of the known sufficient conditions for the RNIEP. It was not until 2017 that it was shown in [2] that C-realizability was more closely related to the Symmetric Nonnegative Inverse Eigenvalue problem (SNIEP)) than to the RNIEP. The combinatorial nature of the original definition of C-realizability has conduced over the years to combinatorial characterizations of the set of C-realizable lists. In this talk we first review a partial one (see [4]) for real lists with zero sum, and then the most general combinatorial characterization, obtained in [5] for arbitrary real lists. One of the most remarkable consequences of the latter is that it proves the monotonicity of C-realizability, i.e., that the operation of increasing any positive entry of a C-realizable list preserves C-realizability.

References

[1] A. Borobia, J. Moro & R. L. Soto (2008). A unified view on compensation criteria in the real nonnegative inverse eigenvalue problem. Linear Algebra Appl., 428, 2574-2584.

[2] R. Ellard & H. Smigoc (2017). Connecting sufficient conditions for the Symmetric Nonnegative Inverse Eigenvalue Problem. Linear Algebra Appl, 498, 521-552.

[3] C. R. Johnson, T. J. Laffey & R. Loewy (1996). The Real and the Symmetric Nonnegative Inverse Eigenvalue Problem are different. Proc. Am. Math. Soc., 124, 3647-3651.

[4] C. Marijuán & J. Moro. (2021) A characterization of trace-zero sets realizable by compensation in the SNIEP. Linear Algebra Appl., 615, 42-76.

[5] C. Marijuán & J. Moro. (2024) A characterization of sets realizable by compensation in the SNIEP. Linear Algebra Appl., 693, 425–447.

NIEP: positive and irreducible realizations

Author: Miriam Pisonero[1]

[1] *Universidad de Valladolid, Spain*

Contact: mpisonero@uva.es

The Nonnegative Inverse Eigenvalue Problem (NIEP) is the problem of characterizing the lists σ of n complex numbers (counting multiplicities) that are the spectrum of a nonnegative matrix of size n. A list σ is said to be realizable if there exists a nonnegative matrix whose spectrum is σ.

Another way of facing the NIEP is to focus the attention on the characteristic polynomial $P(x)$ of σ, and to characterize the real monic polynomials of degree n that are the spectrum of a nonnegative matrix of size n. Note that a nonnegative matrix A can be seen as the adjacency matrix of a weighted digraph. A real monic polynomial $P(x)$ is said to be realizable if there is a weighted digraph, or equivalently a nonnegative matrix, whose characteristic polynomial is $P(x)$.

Our interest here lies in identifying, among realizable spectra/polynomials, those that are realizable by an irreducible matrix/a strongly connected weighted digraph. We use a special type of digraph structure, called EBL, in which the cycles are especially simple. The EBL digraphs are a tool that allow to relate the information contained in the cyclic structure of the digraph associated to A to the coefficients of its characteristic polynomial $P(x)$, and vice versa, in a suitable way.

After giving some general background, we make some useful new observations and show the existence of irreducible nonnegative and positive realizations in some general cases. We discuss both types of realizations for $n < 5$, where the NIEP is solved. Finally, we study the trace 0 case and, using graph theoretic methods and EBL digraphs, characterize nonnegative irreducible realizability among realizable polynomials.

This talk is based on a joint work with C.R. Johnson and C. Marijuán.

Characteristic Polynomials of Nonnegative Hessenberg Matrices

Author: Helena Šmigoc[1]

[1] *University College Dublin*

Contact: helena.smigoc@ucd.ie

The nonnegative inverse eigenvalue problem (NIEP) seeks to characterize the multisets of complex numbers that can be realized as the spectra of nonnegative matrices. Hessenberg matrices have been used in several central constructive results within this field, including the resolution of the NIEP for 4×4 matrices, the characterization of spectra where all non-Perron eigenvalues possess negative real parts, and the constructive proof of the Boyle-Handelman theorem.

In this talk, we examine the set of characteristic polynomials realized by nonnegative Hessenberg matrices. To analyze the underlying zero-nonzero patterns of this family of matrices, we introduce an adaptation of the non-symmetric strong spectral property. We apply this modified spectral property to address the following problem: given a prescribed pattern of a nonnegative Hessenberg matrix A, we seek to identify perturbations of the characteristic polynomial of A that maintain its realizability with nonnegative Hessenberg matrices.

Eigenvalue regions and realising matrices for monotone stochastic matrices

Authors: Brando Vagenende[1], Brecht Verbeken[1], Marie-Anne Guerry[1]

[1] *Department Business Technology and Operations, Data Analytics Laboratory, Vrije Universiteit Brussel (VUB)*

Contact: brando.vagenende@vub.be

This talk presents spectral properties of monotone stochastic matrices which are characterised by the fact that each row stochastically dominates the preceding one, and which arise in contexts such as intergenerational mobility, equal-input models, and credit-rating systems.

In analogy with the stochastic matrices, for the monotone stochastic matrices both the individual eigenvalues as the spectrum as a whole are examined. Individually, the eigenvalue region for all $n \times n$ monotone matrices with $1 \leq n \leq 3$ is completely determined, and realising matrices are provided. Collectively, the set of possible pairs of non-trivial eigenvalues arising from 3×3 monotone matrices is characterised, accompanied by realising matrices. In both perspectives, the resulting regions are substantially smaller than those for general stochastic matrices. Finally, a reduction theorem is presented, stating that, for $n \geq 4$, the eigenvalue region of $n \times n$ monotone matrices is contained within that of $(n-1) \times (n-1)$ stochastic matrices.

Spectrally Complete Subsets and Eigenvalue Regions of classes of Stochastic Matrices

Authors: Brecht Verbeken[1], Vincent Ginis[2]

[1] *Vrije Universiteit Brussel*

[2] *Vrije Universiteit Brussel, Harvard University*

Contact: brecht.verbeken@vub.be

Karpelevich's theorem describes the single-eigenvalue region

$$\Theta_n = \{\lambda \in \mathbf{C} : \lambda \in \sigma(A) \text{ for some } A \in \mathbf{R}^{n\times n} \text{ row-stochastic}\}.$$

The set of row-stochastic matrices is the polytope $\mathrm{conv}(V_n)$, where V_n consists of the n^n deterministic Markov kernels (0-1 matrices with exactly one 1 in each row). For a stochastic class $C = \mathrm{conv}(V)$, call $W \subseteq V$ spectrally complete if

$$\Theta(\mathrm{conv}(W)) = \Theta(C),$$

and minimal if it is inclusion-minimal. Since these eigenvalue regions are typically star-shaped about 0, spectral completeness is essentially radial, writing $\rho_C(\theta) = \sup\{r \geq 0 : re^{i\theta} \in \Theta(C)\}$, equality $\rho_{\mathrm{conv}(W)}(\theta) = \rho_C(\theta)$ for all θ implies $\Theta(\mathrm{conv}(W)) = \Theta(C)$.

I develop this subset-selection viewpoint in three settings.
(1) Doubly stochastic matrices: vertices are permutation matrices; simulations up to $n = 25$ support the Harlev-Johnson-Lim Boundary Conjecture, with no new phenomena beyond the exceptional $n = 5$ behavior, and indicate that a very small family of permutation pairs is already spectrally complete.
(2) Monotone stochastic matrices: building on work of Vagenende et al., I focus on the explicit boundary description for $n = 3$ and explain how AI-assisted was used to identify extremal families and supporting inequalities that were then verified rigorously.
(3) Prescribed zero patterns: in the nonnegative inverse eigenvalue problem with fixed support (Ran-Teng, dimension 3), restricting the zero pattern yields another spectral subset-selection problem.

I conclude by outlining how AlphaEvolve-style workflows could systematize the discovery of spectrally complete subsets and hopefully guide us towards alternative Karpelevich-type proofs.

MS 21: Krylov Iterative Methods for Linear Equations

Minisymposium Organizer: Ron Morgan

Solution of some important problems in linear algebra will be addressed with new iterative approaches. There will not be anything sketchy or mixed, but terms like polynomial, preconditioning, infinite, rank-one, twin and singular may be heard.

List of Talks

Polynomial Preconditioning for Indefinite Matrices

Author: Hayden Henson[1]

[1] *Baylor University*

Contact: hayden_henson1@baylor.edu

Preconditioning plays a central role in accelerating the convergence of iterative methods for solving large linear systems. Among the various approaches, polynomial preconditioning offers a flexible approach. Krylov subspace methods provide a natural setting for constructing polynomials that can be used as preconditioners. In this work, we investigate the use of polynomial preconditioning using the GMRES polynomial on indefinite linear systems. Several challenges can arise in polynomial preconditioning for indefinite matrices. These are mentioned along with some algorithmic strategies designed to mitigate their effects and enhance robustness.

Convergence Analysis for Infinite GMRES

Authors: Jordan Jackson[1], Mark Embree[1]

[1] *Virginia Tech*

Contact: jgjacks4@vt.edu

Many applications require the solution of large-scale linear systems that have nonlinear parameter dependence. The Infinite GMRES algorithm, developed by Jarlebring and Correnty in 2022, converts these systems into infinite-dimensional systems with linear parameter dependence. This transformation involves a linearization process that results in a system with a special companion-like structure. Due to the shift invariance of the Krylov subspaces, all choices of parameter give the same Krylov subspaces. Thus performing the Arnoldi process for one parameter enables the efficient construction of the GMRES iterates for all parameter values. The special structure of the right-hand-side vector of the infinite-dimensional system allows this Arnoldi process to be executed with finite-dimensional vectors, where the dimension grows with every iteration. As such, Infinite GMRES uses finite-length Krylov basis vectors to solve an infinite-dimensional system. The special structure of the infinite-dimensional system invites further investigation into the convergence behavior of the Infinite GMRES residual norm. We aim to understand how various components of this system, including the size and nature of the block matrices, the number of blocks, and the choice of parameter value influence convergence behavior. In this presentation, we will provide rigorous convergence results using various tools of analysis, including the numerical range.

Polynomial Preconditioned Golub-Kahan for Finding Singular Values and Solving Least Squares

Authors: Hayden Henson[1], Kingsley Michael[1], Nathaniel Morgan[1], Ron Morgan[1]

[1] *Baylor University*

Contact: kingsley_michael1@baylor.edu

Golub-Kahan bidiagonalization procedure is well known for its role in computing singular values and solving least squares problems. We present a polynomial-preconditioned variant of this framework aimed at reducing the need for restarting and extensive orthogonalization. This talk outlines the formulation and motivation for the method and examines its potential to improve performance in practical settings.

Krylov Methods for Rank-one Updates of Eigenvalue Problems and Linear Equations

Author: Dona Ishara Saparamadu[1]

[1] *Baylor University, Waco, Texas*

Contact: ishara_saparamadu1@baylor.edu

Krylov methods are given for rank-one updates of both eigenvalue and linear equations problems. For eigenvalues, an Arnoldi iteration for the original matrix can be continued on the rank-one changed matrix. We discuss how careful implementation allows the desired accuracy to be attained for the updated matrix. Next, methods are given for linear equations, one that uses the updated Arnoldi approach and the other has the Sherman-Morrison formula. We also give some theoretical results about how eigenvalues change with rank-one perturbation.

Preconditioner Updating with Lasso-based Sparse Approximate Maps

Authors: Achintya Sunil, Eric De Sturler
Contact: axs2061@vt.edu

Design problems such as topology optimization and PDE-based inverse problems require the solution of a sequence of linear-systems derived from finite element or finite difference discretization. Preconditioning is essential for the fast solution of these systems by iterative methods. However, computing an accurate preconditioner for every system in the sequence may be a computational bottleneck, warranting a cheap preconditioner-updating scheme. The sparse approximate map (SAM), introduced in Preconditioning Parametrized Linear Systems by Carr et al. (SISC 2021), computes a sparse approximation of the exact map between matrices in the sequence. Each column of the SAM is computed by solving a small least-squares problem. The resulting map can be leveraged to update a preconditioner computed earlier in the sequence. This method is cheap as long as the underlying sparsity pattern has few non-zeros. However, as revealed in prior research, often there is no intuitive best choice for such a pattern. Often the sparsity pattern of the matrix (or powers of it) can be used, but for topology optimization matrices or for high-order finite element stiffness matrices with a large number of non-zeros per column this is a poor choice. In this talk, we show how generating patterns using an L-1 constrained least squares solve (lasso) can be very effective. We solve the lasso problem to modest accuracy using the alternating direction method of multipliers (ADMM). We demonstrate that the parameters of ADMM and the regularization parameter can be heuristically selected to compute matrix maps that are significantly sparse and sufficiently accurate. This is joint work with Johann Rudi and Eric de Sturler.

Twin CG for Linear Equations with Multiple Right-hand Sides

Author: Abigail Williams[1]

[1] *Baylor University*

Contact: abby_williams5@baylor.edu

A very simple approach to solving multiple right-hand side systems is proposed. For symmetric problems, the conjugate gradient method is a very efficient way to solve linear equations. We will use the same parameters from solving the first system for other systems. This is called Twin CG. It corresponds to applying the same polynomial to the other systems as was used for the first system. No dot products are needed except for the first system, and the method is very parallelizable. Deflation of eigenvalues using seeding can be included. The remarkable natural stability control due to roundoff error in the symmetric Lanczos iteration will be discussed. Added stability control methods include deflating eigenvalues with linear factors and applying partial Global CG.

MS 22: Algebraic Invariants of Graphs

Minisymposium Organizers: Carlos A. Alfaro, Ralihe Raul Villagran Olivas

Algebraic graph theory has traditionally focused on the interplay between the combinatorial properties of graphs and the linear algebraic properties of their associated matrices. Recently, the field has seen advances in the study of certain algebraic structures of the distance matrices of graphs, Laplacian matrices of graphs, and variants. In particular, a focus was placed on the spectrum, Smith normal form, arithmetical structures, sandpile groups, and determinantal ideals, among others. The purpose of this session is to create a diverse forum for students and researchers interested in algebraic graph theory and its associated algebraic invariants. The session will highlight how these algebraic invariants can distinguish graphs that are cospectral with respect to standard (adjacency and Laplacian) matrices, and how they relate to combinatorial parameters such as the domination number, diameter, and zero forcing number. The talks will bridge the gap between Combinatorial Matrix Theory, Group Theory, and Commutative Algebra, in the context of Linear Algebra.

List of Talks

The Smith normal form of distance matrices of high dimensional trees

Authors: Carlos Alfaro[1], Iván Téllez Téllez, Jesús Uriel Medrano

[1] *Banxico*

Contact: alfaromontufar@gmail.com

Graham, Lovász and Pollak obtained a well known formula for the determinant of distance matrices of trees. This formula depends only on the number of vertices of the tree and not on its topological

structure. Later, Hou and Woo computed the explicit expression of the Smith form of the distance matrix of a tree, which again, it depends only on the number of vertices. In this talk, we will show results that extend Graham-Lovász-Pollak and Hou-Woo formulas to higher dimensions by computing the Smith normal form and the determinant of a distance matrix associated to high dimensional trees.

Connected domination and the v-number of binomial edge ideals

Authors: Delio Jaramillo Velez[1], Lisa Seccia[2]

[1] *Virginia Tech*

[2] *University of Neuchâtel*

Contact: delio@chalmers.se

A connected dominating set of a graph is a vertex set that induces a connected subgraph and such that every vertex outside the set is adjacent to at least one vertex in the set. The minimum cardinality of a connected dominating set is called the connected domination number. We present an algebraic expression for this combinatorial invariant using the theory of binomial edge ideals.

Spectrum of trees of diameter 4 for the distance Laplacian matrix

Author: Kate Lorenzen[1]

[1] *Linfield University*

Contact: klorenzen@linfield.edu

Graphs can be encoded into a matrix according to some rule. The eigenvalues of the matrix are used to understand the structural properties of graphs. If two graphs share a set of eigenvalues, they are called cospectral. A tree is a graph with no cycles, and for most matrix representations, almost all trees have a cospectral mate. The distance Laplacian matrix is found by subtracting the distance matrix from the diagonal transmission matrix. There is an open conjecture that trees, for their distance Laplacian matrices, do not have a cospectral mate and are therefore spectrally determined. We explicitly find the spectrum of a family of trees of diameter 4 and how it is uniquely defined by its parameters.

Generalized chip firing and critical groups of arithmetical structures on trees

Authors: Kassie Archer[1], Alexander Diaz-Lopez[2], Darren Glass[3], Joel Louwsma[4]

[1] United States Naval Academy

[2] Villanova University

[3] Dickinson College

[4] Niagara University

Contact: jlouwsma@niagara.edu

Chip firing provides a way to study the sandpile group (also known as the Jacobian) of a graph. We use a generalized version of chip firing to bound the number of invariant factors of the critical group of an arithmetical structure on a graph. We also show that, under suitable hypotheses, critical groups are additive under wedge sums of graphs with arithmetical structures. These results allow us to relate the number of invariant factors of critical groups associated to any given tree to decompositions of the tree into simpler trees. We use this to classify those trees for which every arithmetical structure has cyclic critical group. Finally, we show how to construct arithmetical structures on trees with prescribed critical groups. In particular, every finite abelian group is realized as the critical group of some arithmetical structure on a tree.

Smith Normal Forms of Graphical Hermite Simplices

Authors: Antwon Park[1], Benjamin Braun[1]

[1] University of Kentucky

Contact: antwon.park@uky.edu

We introduce the family of graphical Hermite simplices and study the Smith normal forms of their matrices of vertex vectors, which is equivalent to studying the group structure of the cokernels for these matrices. Our motivation is to study the behavior of lattice simplices subject to small lattice perturbations of their vertices. In this case, a graphical Hermite simplex is a perturbation of a rectangular simplex, i.e., a simplex defined by a diagonal matrix and the origin, with the perturbation controlled by the structure of a directed graph. We first establish sufficient conditions on the graphs and diagonal entries of these matrices that imply having a single non-unit invariant factor, i.e., a cyclic cokernel. We then obtain bounds on the invariant factors of the defining matrices related to lengths of paths in the corresponding directed graph.

The characterization of graphs with two trivial distance ideals

Authors: Carlos Alfaro[1], Juan Pablo Serrano Perez[2], Ralihe Villagran[3], Teresa Hoekstra Mendoza[4]

[1] Banco de Mexico

[2] Cinvestav

[3] *WIP*

[4] *Cimat*

Contact: wamjsblvbnp@gmail.com

The distance ideals of connected graphs are algebraic invariants extending the Smith normal form (SNF) and the spectrum of graph distance matrices. In general, distance ideals are not monotone under taking induced subgraphs. However, it was proved in 2017 that the set of graphs with one trivial distance ideal over $\mathbb{Z}[X]$ and over $\mathbb{Q}[X]$ was characterized in terms of induced subgraphs, where X is a set of variables indexed by the vertices. Here, we give a characterization of the $\{\mathcal{F}, \mathsf{odd\text{-}holes}_7\}$-free graphs, where $\mathsf{odd\text{-}holes}_7$ consists of the odd cycles of length at least seven and $\mathcal{F}$ is a set of sixteen graphs. Moreover, we show that the $\{\mathcal{F}, \mathsf{odd\text{-}holes}_7\}$-free graphs are precisely the graphs with at most two trivial distance ideals over $\mathbb{Z}[X]$. As byproduct, we also find that the determinant of the distance matrix of a connected bipartite graph is even, this suggests that it is possible to extend, to bipartite graphs, the Graham-Pollak-Lovász celebrated formula $\det(D(T_{n+1})) = (-1)^n n 2^{n-1}$, and Hou-Woo's result stating that $\mathrm{SNF}(D(T_{n+1})) = \mathsf{I}_2 \oplus 2\mathsf{I}_{n-2} \oplus (2n)$, for any tree T_{n+1} with $n+1$ vertices. Furthermore, we also determine the graphs with at most two trivial distance ideals over $\mathbb{Q}[X]$, and the graphs with at most two trivial distance univariate ideals. We conclude by showing that the SNF of the distance matrix of a graph has exactly two invariant factors equal to 1 if and only if it is bipartite or complete tripartite.

A representation theoretical approach to the p-rank of subset incidence matrices

Authors: Colby Sherwood[1], Joshua Ducey

[1] *University of Delaware*

Contact: sherwocj@udel.edu

Let $W^i_{k,n}(m)$ denote a matrix with rows and columns indexed by the k-subsets and n-subsets, respectively, of an m-element set. The row S, column T entry of $W^i_{k,n}(m)$ is 1 if $|S \cap T| = i$, and is 0 otherwise. When $i = k$ the matrix $W^k_{k,n}(m)$ is the subset inclusion matrix for which Wilson found a diagonal form, solving the p-rank problem for any prime p. This diagonal form was used to calculate the Smith group of the hypercube graph.

We compute the rank of the matrix $W^1_{2,n}(m)$ over any field by making use of the representation theory of the symmetric group. We also give a simple condition under which $W^i_{k,n}(m)$ has large p-rank.

Advances on Brouwer's Conjecture

Author: Vilmar Trevisan[1]

[1] *UFRGS - Brazil*

Contact: trevisan@mat.ufrgs.br

Within Spectral Graph Theory, Brouwer's Conjecture (BC) is a fundamental problem concerning Laplacian eigenvalues and graph invariants. It proposes a relationship between the sum of the largest Laplacian eigenvalues of a graph and its number of edges, with direct implications for the study of Laplacian energy. More precisely, for a graph $G = (V, E)$ with $n = |V|$ vertices and $m = |E|$ edges, the conjecture states that, for $k = 1 \ldots, n$, the sum $S(k)$ of the k largest eigenvalues of the Laplacian matrix of G satisfies

$$S(k) \leq m + k(k-1)/2.$$

In this talk, we discuss the role of Brouwer's Conjecture in the area of Spectral Graph Theory and describe recent progress on the problem, highlighting results and contributions obtained in this framework.

From Chip Firing to Determinantal Ideals and Back

Author: Ralihe Raul Villagran Olivas[1]

[1] *Worcester Polytechnic Institute*

Contact: rvillagran@wpi.edu

Chip-firing is a discrete dynamical process on a graph that exhibits striking phenomena, including fractal-like symmetries and self-organized criticality. From an algebraic perspective, the states of this process (combinatorially) define a group, the Sandpile group, whose algebraic structure is given by the Smith normal form of the Laplacian matrix. We will discuss problems related to Sandpile Groups and introduce some determinantal ideals for graphs. We will discuss how these algebraic invariants provide a framework for solving said problems in graph theory and, conversely, how the structure of Sandpile groups offers new insights into the properties of these ideals.

MS 23: Low-Complexity Data-driven or Classical Algorithms and Applications

Minisymposium Organizers: Sirani M. Perera, James Nagy, Ilias Kotsireas

The principles of applied linear algebra can be utilized to address numerous scientific and engineering problems. One may start with established techniques but shift towards low-complexity algorithms to find solutions in these areas. Others might investigate the hidden structures that exist within systems, allowing for the design of efficient data-driven techniques that solve challenges in science and engineering. Consequently, we propose leveraging structure, sparsity, symmetry, unitarity, and low-rank approximations to significantly enhance computational efficiency, leading to solving problems in science and engineering. Our focus lies at the forefront of scientific and engineering advancements, bridging the theoretical frameworks of applied linear algebra to meet the demands of complex and high-dimensional data structures. We address how these structures can be utilized to minimize computational demand with accuracy and stability in solving linear or nonlinear systems. Thus, the symposium will explore how the theories of linear algebra can be utilized not only for addressing low-complexity classical algorithms but also for machine learning models that solve pertinent problems in science and engineering. We welcome the ILAS participants to our mini-symposium, where we will explore cutting-edge computational techniques that extend beyond conventional methods. We will address conventional algorithms and machine learning models influenced by linear algebra theories, fostering interdisciplinary collaboration across computational mathematics, scientific computing, data science, engineering, and physics.

List of Talks

Spotlight inversion by orthogonal projections

Authors: Daniela Calvetti[1], Erkki Somersalo[1], Nuutti Hyvonen[2], Ville Kolehmainen[3]

[1] *Case Western Reserve University*

[2] *Aalto University*

[3] *University of Eastern Finland*

Contact: dxc57@case.edu

Many computational problems involve solving a linear system of equations, although only a subset of the entries of the solution are needed. In inverse problems, where the goal is to estimate unknown parameters from indirect noisy observations, it is not uncommon that the forward model linking the observed variables to the unknowns depends on variables that are not of primary interest, often referred to as nuisance parameters. In this talk, we consider linear problems, and propose a novel projection technique to eliminate, or at least mitigate, the contribution of the nuisance parameters in the model. We refer to this approach as spotlight inversion, as it allows to focus on only the portion of primary interest of the unknown parameter vector, leaving the uninteresting part in the shadow. The viability of the approach is illustrated with two computed examples.

Novel Randomized Tensor-Train Sketch and Applications

Authors: Paul Cazeaux[1], Mi-Song Dupuy[2], Rodrigo Figueroa Justiniano[1]

[1] *Virginia Tech*

[2] *Sorbonne University*

Contact: cazeaux@vt.edu

The Tensor-Train (TT) or Matrix-Product States (MPS) format provides a compact, low-rank representation for high-dimensional tensors, widely used in many-body quantum physics and quantum chemistry. Its efficiency relies on rounding, which reduces tensor ranks to maintain feasible computational costs. In this talk, we introduce a novel block-structured randomized sketch exploiting the TT format and provide explicit probabilistic guarantees. We discuss how this sketch enables in particular randomized rounding algorithms for TT rounding, extending low-rank matrix approximation techniques. These methods significantly accelerate computations, particularly for summing TT-tensors or performing matrix-vector products, key operations in Krylov iterative solvers. We present numerical results demonstrating the empirical accuracy and computational advantages of the randomized approach over deterministic methods.

This research was supported in part by Simons Travel Support for Mathematicians Award No. MPS-TSM-00966604.

Optimal Experimental Design for Gaussian Processes via Column Subset Selection

Authors: Jessie Chen[1], Arvind Krishna Saibaba[1], Hangjie Ji[1]

[1] *North Carolina State University*

Contact: jchen228@ncsu.edu

Gaussian process regression uses data measured at sensor locations to reconstruct a spatially dependent function with quantified uncertainty. However, if only a limited number of sensors can be deployed, it is important to determine how to optimally place the sensors to minimize a measure of the uncertainty in the reconstruction. We consider the Bayesian D-optimal criterion to determine the optimal sensor locations by choosing k sensors from a candidate set of n sensors. Since this is an NP-hard problem, our approach models sensor placement as a column subset selection problem (CSSP) on the covariance matrix, computed using the kernel function on the candidate sensor points. We propose an algorithm that uses the Golub-Klema-Stewart framework (GKS) to select sensors and provide an analysis of lower bounds on the D-optimality of these sensor placements. To reduce the computational cost in the GKS step, we propose and analyze algorithms for the D-optimal sensor placements using Nyström approximations on the covariance matrix. Moreover, we propose several algorithms that select sensors via Nyström approximation of the covariance matrix, utilizing the randomized Nyström approximation, random pivoted Cholesky and greedy pivoted Cholesky. We demonstrate the performance of our method on two applications: thin liquid film dynamics and sea surface temperature.

A low-complexity locally recoverable matrix-influenced algorithm to globally recover codes

Authors: Oshani Jayawardane[1], Sirani M. Perera[1], Hansaka Aluvihare[2], Xianqi Li[3]

[1] *Embry-Riddle Aeronautical University*

[2] *University of Sri Jayewardenepura*

[3] *Florida Institute of Technology*

Contact: oshani.jayawardane@gmail.com

Code Recovery using algebraic-geometric approaches becomes computationally expensive with the cardinality of the field and the complexity of the code structures. In response, we present a low-complexity algorithm that utilizes structures in algebraic-geometric codes over finite fields. The

low-complexity algorithm recovers algebraic codes over finite fields locally, which we name as ***lrc*** algorithm. The ***lrc*** algorithm is derived based on a sparse matrix factorization of the inverse of an $(r+1) \times r$ locally recoverable matrix, whose elements are defined by the points on the surface in $\mathbb{P}^3$ over the finite field $\mathbb{F}_q$ having locality r. We will show that the ***lrc*** algorithm reduces the complexity from $\mathcal{O}(n^3)$ to $\mathcal{O}(n \log n)$ for the $n = 2^s (s \geq 1) > r$ length codeword. Next, we utilize an extended ***lrc*** algorithm followed by a structured neural network (StNN) to globally recover codes over the surface $\mathbb{P}^3$ over the finite field $\mathbb{F}_q$. We discuss two StNNs, i.e., DFT-StNN and DCT-StNN, based on the factorization of the locally recoverable matrix to globally recover codes. Numerical simulations will be shown to compare the performance of the extended ***lrc*** algorithm in global recovery codes, with brute-force calculation, DFT-StNN, DCT-StNN, and a feedforward neural network for codewords with lengths from $n = 6, 12, 27, 48, 96$ and 210, having locality 2 for points on the surface $\mathbb{P}^3$ over the finite field $\mathbb{F}_q$. Our empirical results demonstrate that the extended ***lrc*** algorithm achieves the lowest flops, the highest accuracy with an error order 10^{-16} compared to brute-force calculations, DFT-StNN, DCT-StNN, and the feedforward neural network, showing the existence of local to global codeword setting.

Conditional Denoising Diffusion Model-Based Robust MR Image Reconstruction from Highly Undersampled Data

Authors: Xianqi Li[1], Mohammed Alsubaie, Sirani Perera, Wenxi Liu

[1] *Florida Institute of Technology*

Contact: xli@fit.edu

Magnetic Resonance Imaging (MRI) is a critical tool in modern medical diagnostics, yet its prolonged acquisition time remains a critical limitation, especially in time-sensitive clinical scenarios. While undersampling strategies can accelerate image acquisition, they often result in image artifacts and degraded quality. Recent diffusion models have shown promise for reconstructing high-fidelity images from undersampled data by learning powerful image priors; however, most existing approaches either (i) rely on unsupervised score functions without paired supervision or (ii) apply data consistency only as a post-processing step. In this work, we introduce a conditional denoising diffusion framework with iterative data consistency correction, which differs from prior methods by embedding the measurement model directly into every reverse diffusion step and training the model on paired undersampled–ground truth data. This hybrid design bridges generative flexibility with explicit enforcement of MRI physics. Experiments on the fastMRI dataset demonstrate that our framework consistently outperforms recent state-of-the-art deep learning and diffusion-based methods in SSIM, PSNR, and LPIPS, with LPIPS capturing perceptual improvements more faithfully. Specifically, under an acceleration factor of 8 and Gaussian 1D sampling, the proposed model achieves SSIM = 0.834 ± 0.063, PSNR = 32.52 ± 2.63 dB, and LPIPS = 0.063 ± 0.029. These results demonstrate that integrating conditional supervision with iterative consistency updates yields substantial improvements in both pixel-level fidelity and perceptual realism, establishing a

principled and practical advance toward robust, accelerated MRI reconstruction.

Randomized Sketching for Tucker Tensors: From Compressed Summation to GMRES

Authors: Mirjeta Pasha, Alberto Bucci[1], Andrés Galindo-Olarte[2], Grey Ballard[3], Hussam Al-Daas[4], Jingmei Qiu[5], Joseph Nakao[6], Martina Iannacito[7], Rudi Smith[8], William Taitano[9]

[1] *University of Edinburgh*

[2] *University of Texas at Austin*

[3] *Wake Forest University*

[4] *Scientific Computing Department, STFC, Rutherford Appleton Laboratory, Didcot OX11 0QX, UK*

[5] *University of Delaware*

[6] *Swarthmore College*

[7] *University of Bologna*

[8] *Virginia Tech*

[9] *Los Alamos National Laboratory*

Contact: mpasha@vt.edu

Tensor equations of the form $L(X) = B$, with L a sum of Kronecker products, arise across scientific computing, from high-dimensional PDEs to large-scale inverse problems. When the tensor order is moderate but mode sizes are very large, the Tucker format is an attractive choice — yet using it inside iterative solvers is notoriously hard, as operator applications, orthogonalization, and linear combinations all trigger rank growth and force expensive truncations of large intermediate tensors.

In this talk, I will present a line of work that addresses these bottlenecks through randomized sketching applied directly to Tucker factors. The first part introduces sketching-based methods for Tucker tensor summation that exploit the structure of Khatri–Rao and Kronecker products to perform compressed arithmetic without ever forming dense intermediates. Building on this foundation, the second part develops a randomized Tucker-sketched GMRES solver in two variants: a robust method based on the randomized HOSVD with adaptive accuracy control, and a faster one built on the multilinear Nyström approximation, which — at the price of a prescribed maximal rank — accelerates orthogonalization, avoids storing the Krylov basis, and assembles the solution at negligible cost. Numerical experiments on image deblurring, a parameter-dependent elliptic problem, a transport equation, and a convection–diffusion problem illustrate the efficiency and robustness of the framework.

Computing Singular Values Above a Certain Threshold

Authors: James Baglama[1], Jonathan Chavez-Casillas[1], Vasilije Perovic[1]

[1] *University of Rhode Island*

Contact: perovic@uri.edu

In this talk we discuss various computational aspects of determining all singular triplets $\{\sigma_j, u_j, v_j\}$ corresponding to singular values of A above a user-specified threshold parameter *sigma*, or in other words, determining a k-PSVD of A such that $\sigma_k \geq sigma$ and $\sigma_{k+1} < sigma$. While various numerical schemes with publicly available software have been developed for finding an approximate k–PSVD of a large sparse matrix, e.g., PROPACK, `irlba`, `primme_svds`, `rsvd`, and MATLAB's `svds`, their use is somewhat limited in this context as they all require knowing the number of desired singular triplets, k, in advance.

We propose a hybrid algorithm that works directly on A and is based on the *explicit deflation procedure applied to Golub-Kahan-Lanczos Bidiagonalization (GKLB)-based methods*. An important feature of GKLB-methods that is relevant here is that they can be thought of as a one–sided PSVD approximations of A – this allows for the explicit deflation technique to be applied only to either the left or right singular vectors of A, depending on its dimensions, and thus significantly reducing the overall computational cost. Due to iterative deflation process, a modest non–zero error is added to the side in the GKLB method that had theoretically zero error. Although this error growth is theoretically manageable, our numerical experiments showed that at times one has to restore the one–sided GKLB structure which we have achieved by combining the GKLB-based method with a variant of the block SVD power method. Implementations of this algorithm are provided in MATLAB and a widely used statistical software R along with alternative thresholding criteria based on the notions of *energy* and the *normalized root mean squared error*.

Convergence of Randomized and Greedy Block Gauss-Seidel methods, as well as Asynchronous Iterations

Authors: Daniel Szyld[1], Andreas Frommer[2]

[1] *Temple University*

[2] *University of Wuppertal*

Contact: szyld@temple.edu

We extend results known for the randomized (point and block) Gauss-Seidel and the Gauss-Southwell methods for the case of a Hermitian and positive definite matrix to certain classes of non-Hermitian matrices. We consider cases with overlapping variables (as in Domain Decomposition). We obtain convergence results for a whole range of parameters describing the probabilities in the randomized method or the greedy choice strategy in the Gauss-Southwell-type methods. We identify those choices which make our convergence bounds best possible. One result is that the best convergence bounds that we obtain for the expected values in the randomized algorithm are

as good as the best for the deterministic, but more costly algorithms of Gauss-Southwell type. We use these new results to show a provable convergence rate for asynchronous iterations.

Accelerating Spectral Clustering of Time Series by approximating the Similarity Matrix using Randomly Pivoted Cholesky

Authors: Raf Vandebril[1], Sander Das[2], Thijs Steel[2], Wannes Meert

[1] *Dept. Computer Science, KU Leuven*

[2] *KU Leuven*

Contact: raf.vandebril@kuleuven.be

A major problem for time series clustering is that computing the similarity matrix for the most used similarity measures becomes infeasible if number amount or length of time series becomes too large. However, since this similarity matrix typically has low-rank structure, it can be approximated using a low-rank approximation. In this work, we show that existing numerical linear algebra methods, more specifically Randomly Pivoted Cholesky can be used in the context of time series clustering to drastically reduce the computational cost of calculating the similarity matrix, while maintaining the clustering quality. This shows that low-rank approximation algorithms are an effective and scalable technique that can be used in time series clustering.

MS 24: Approximate Computing in Numerical Linear Algebra

Minisymposium Organizers: Massimiliano Fasi, Xiaobo Liu

Approximate computing techniques have gained substantial attention over the last two decades, driven by the growing need for scalable and energy-efficient algorithms in scientific computing and data-intensive applications. In numerical linear algebra, they have become increasingly prevalent, and todaly they enable significant perormance gains without sacrificing accuracy. Methods such as randomization, low-rank approximation, and mixed-precision arithmetic have all shown their potential to accelerate computations, reduce memory footprints, and lower energy consumption, extending the range of problems that can be solved on modern computing systems. The impact of approximate computing is particularly evident in large-scale simulations and machine learning, where traditional exact algorithms often become prohibitively expensive. Recent developments open new opportunities for designing algorithms that balance accuracy and robustness without sacrificing performance. This minisymposium aims to bring together researchers working on these topics to share recent breakthroughs and discuss emerging trends. Contributions will explore approximate computing in the context of a broad spectrum of numerical linear algebra kernels, including—but not limited to—solvers for linear systems and least squares problems, eigenvalue computations, matrix functions, and matrix equations.

List of Talks

Automated Precision Tuning for Numerical Algorithms

Authors: Xinye Chen[1], Thibault Hilaire[2], Fabienne Jézéquel[3]

[1] *LIP6, CNRS, Sorbonne University*

[2] *LIP6, Sorbonne University*

[3] *Sorbonne Université, CNRS, LIP6 Université Paris-Panthéon-Assas*

Contact: xinyechenai@gmail.com

Achieving optimal performance in numerical computations often hinges on aggressively reducing precision or performing rigorous rounding-error analysis to retain numerical accuracy. The precision tuning methods provide a unified, task-specific validation platform for automated precision tuning, enabling a balance between computational efficiency and numerical fidelity. In this talk, we present novel precision-tuning methods, including expanded precision options and low-precision math functions, which streamline the exploration of low-precision configurations. This talk reveals the potential for significant performance gains by leveraging low-precision settings. This work underscores the transformative potential of automated precision tuning while maintaining accuracy and performance objectives.

A mixed precision algorithm for matrix root functions

Authors: Bowen Gao[1], Daniel Kressner[2], Meiyue Shao[1]

[1] *Fudan University*

[2] *EPFL*

Contact: bwgao22@m.fudan.edu.cn

Mixed precision computation has attracted great attention in recent years partly due to the evolution of machine learning and hardware infrastructure. Recent development on mixed precision algorithms has largely enhanced the performance of various linear algebra solvers. In this talk, we propose a mixed precision algorithm for the computation of matrix root functions, primarily the matrix square root and other higher order roots. We introduce a new refinement framework that is capable of refining the lower precision solution to the working precision level. We show that carefully designed mixed precision algorithms compute the matrix root functions to full working precision accuracy and offer speedup in most scenarios compared to the fixed precision method.

Adaptive Mixed Precision Hierarchical Matrices

Authors: Ritesh Khan, Erin Carson
Contact: ritesh.khan@matfyz.cuni.cz

Hierarchical matrices or $\mathcal{H}$-matrices are the block low-rank representation of the original matrices and are widely used in fast matrix computations. In this talk, we show that the low-rank blocks of $\mathcal{H}$-matrices can be represented in low precision (precision lower than the working precision) without degrading the overall approximation quality. We provide an explicit rule to decide which precision should be chosen for a particular low-rank block. We propose an adaptive mixed precision algorithm for constructing and storing $\mathcal{H}$-matrices. We also show that use of mixed precision does not compromise the numerical stability and accuracy of the $\mathcal{H}$-matrix-vector product. We perform a wide range of numerical experiments to validate our theoretical results. Our numerical results illustrate that adaptive mixed precision $\mathcal{H}$-matrices can achieve significant storage reductions compared to uniform precision $\mathcal{H}$-matrices, without compromising accuracy.

A general framework for Krylov ODE residuals with applications to randomized Krylov methods

Authors: Emil Krieger[1], Marcel Schweitzer[1]

[1] *Bergische Universität Wuppertal*

Contact: krieger@uni-wuppertal.de

Randomized Krylov subspace methods that employ the sketch-and-solve paradigm to substantially reduce orthogonalization cost have recently shown great promise in speeding up computations for many core linear algebra tasks (e.g., solving linear systems, eigenvalue problems and matrix equations, as well as approximating the action of matrix functions on vectors) whenever a non-symmetric matrix is involved. An important application that requires approximating the action of matrix functions on vectors is the implementation of exponential integration schemes for ordinary differential equations. In this paper, we specifically analyze randomized Krylov methods from this point of view. In particular, we use the residual of the underlying differential equation to derive a new, reliable a posteriori error estimate that can be used to monitor convergence and decide when to stop the iteration. To do so, we first develop a very general framework for Krylov ODE residuals that unifies existing results, simplifies their derivation and allows extending the concept to a wide variety of methods beyond randomized Arnoldi (e.g., rational Krylov methods, Krylov methods using a non-standard inner product, ...). In addition, we discuss certain aspects regarding the efficient implementation of sketched Krylov methods. Numerical experiments on large-scale ODE models from real-world applications illustrate the quality of the error estimate as well as the general competitiveness of sketched Krylov methods for ODEs in comparison to other Krylov-based methods.

Forward and backward error bounds for a mixed precision preconditioned conjugate gradient algorithm

Author: Yuxin Ma[1]

[1] *Charles University*

Contact: yuxin.ma@matfyz.cuni.cz

The preconditioned conjugate gradient (PCG) algorithm is one of the most popular algorithms for solving large-scale linear systems $Ax = b$, where A is a symmetric positive definite matrix. Rather than computing residuals directly, it updates the residual vectors recursively. Current analyses of the conjugate gradient (CG) algorithm in finite precision typically assume that the norm of the recursively updated residual goes orders of magnitude below the machine precision, focusing mainly on bounding the residual gap thereafter. This work introduces a framework for the PCG algorithm and provides rigorous proofs that the relative backward and forward errors of the computed results of PCG can reach the levels $O(u)$ and $O(u)\kappa(A)^{1/2}$, respectively, after a sufficient number of iterations without relying on an assumption concerning the norm of the recursively updated residual, where u represents the unit roundoff and $\kappa(A)$ is the condition number of A. Our PCG framework further shows that applying preconditioners in low precision does not compromise the accuracy of the final results, provided that reasonable conditions are satisfied. Our theoretical results are illustrated through a set of numerical experiments.

Restoring similarity in randomized Krylov methods with applications to eigenvalue problems and matrix functions

Authors: Laura Grigori[1], Daniel Kressner[1], Nian Shao[1], Igor Simunec[1]

[1] *EPFL*

Contact: igor.simunec@epfl.ch

The randomized Arnoldi process has been used in large-scale scientific computing because it produces a well-conditioned basis for the Krylov subspace more quickly than the standard Arnoldi process. However, the resulting Hessenberg matrix is generally not similar to the one produced by the standard Arnoldi process, which can lead to delays or spike-like irregularities in convergence. In this talk we introduce a modification of the randomized Arnoldi process that restores similarity with the Hessenberg matrix generated by the standard Arnoldi process. This is accomplished by enforcing orthogonality between the last Arnoldi vector and the previously generated subspace, which requires solving only one additional least-squares problem. When applied to eigenvalue problems and matrix function evaluations, the modified randomized Arnoldi process produces approximations that are identical to those obtained with the standard Arnoldi process. Numerical experiments demonstrate that our approach is as fast as the randomized Arnoldi process and as robust as the standard Arnoldi process.

This talk is based on a joint work with Laura Grigori, Daniel Kressner and Nian Shao (https://arxiv.org/abs/2601.10248).

Analysis of Floating-Point Matrix Multiplication Computed via Integer Arithmetic

Authors: Ahmad Abdelfattah[1], Jack Dongarra[1], Massimiliano Fasi[2], Mantas Mikaitis[2], Francoise Tisseur[3]

[1] *University of Tennessee*

[2] *University of Leeds*

[3] *The University of Manchester*

Contact: massimiliano.fasi@gmail.com

Ootomo, Ozaki, and Yokota [*Int. J. High Perform. Comput. Appl.*, 38 (2024), p. 297–313] have proposed a strategy to recast a floating-point matrix multiplication in terms of integer matrix products. The factors A and B are split into integer *slices*, the product of these slices is computed exactly, and AB is approximated by accumulating these integer products in floating-point arithmetic. This technique is particularly well suited to mixed-precision matrix multiply–accumulate units with integer support, such as the NVIDIA tensor cores or the AMD matrix cores The number of slices allows for performance-accuracy tradeoffs: more slices yield better accuracy but require more multiplications, which in turn reduce performance. We propose an inexpensive way to estimate the minimum number of multiplications needed to achieve a prescribed level of accuracy. Our error analysis shows that the algorithm may become inaccurate (or inefficient) if rows of A or columns of B are *badly scaled*. We perform a range of numerical experiments, both in simulation and on the latest NVIDIA GPUs, that confirm the analysis and illustrate strengths and weaknesses of the algorithm.

Iterative Refinement for a Subset of Eigenpairs of a Real Symmetric Matrix and Its Convergence Analysis

Authors: Takeshi Terao[1], Katsuhisa Ozaki[2], Toshiyuki Imamura[3]

[1] *Waseda University*

[2] *Shibaura Institute of Technology*

[3] *RIKEN Center for Computational Science*

Contact: takeshi.terao@aoni.waseda.jp

We consider the eigenvalue problem $Ax^{(i)} = \lambda_i x^{(i)}$ for a real symmetric matrix $A \in \mathbb{R}^{n\times n}$, where $\lambda_i \in \mathbb{R}$ is an eigenvalue of A and $x^{(i)} \in \mathbb{R}^n$ is the corresponding eigenvector. This work investigates iterative refinement methods to improve the accuracy of eigenvectors $x^{(i)}$.

Efficient methods are known for improving the accuracy of either a single approximate eigenvector or all approximate eigenvectors simultaneously. In practical applications, however, it is often necessary to improve only a subset of eigenvectors: $X = \left(x^{(p_1)}, x^{(p_2)}, \ldots, x^{(p_k)}\right) \in \mathbb{R}^{n\times k}, \{p_i\} = \{1, \ldots, n\}, 1 \le k \le n$.

The proposed method improves the accuracy of an approximation $\widehat{X}$ to X. Given the matrix A and an approximate eigenvector matrix $\widehat{X}$, the goal is to compute $\widetilde{X}$ such that $\|X - \widetilde{X}\| < \|X - \widehat{X}\|$. A necessary condition for improving accuracy is $\min_{1\leq i\leq k} |\lambda_{p_i}| > \max_{k+1\leq j\leq n} |\lambda_{p_j}|$.

The proposed method has two main advantages: its computational cost is dominated by matrix multiplications, and it can handle large sparse matrices, making it applicable to mixed-precision and high-precision numerical computations. We report both analytical results and numerical experiments.

Mixed Precision General Alternating-Direction Implicit Method for Solving Large Sparse Linear Systems

Authors: Bastien Vieublé, Jifeng Ge, Juan Zhang
Contact: zhangjuan@xtu.edu.cn

Our work presents a novel mixed precision formulation of the General Alternating-Direction Implicit (GADI) method, designed to accelerate the solution of large-scale sparse linear systems $Ax = b$. By solving the computationally intensive subsystems in low precision (e.g., Bfloat16 or FP32) and performing residual and solution updates in high precision, the proposed method significantly reduces execution time without compromising the final solution accuracy. We provide a comprehensive rounding error analysis that establishes convergence rates and limiting accuracies under mixed precision arithmetic. Additionally, we introduce a robust strategy based on Gaussian Process Regression (GPR) for selecting the optimal regularization parameter α. Performance benchmarks on an NVIDIA A100 GPU demonstrate that mixed precision GADI achieves speedups of up to 3.1× compared to standard double precision implementations on large-scale convection-diffusion and reaction-diffusion problems.

MS 25: Rational Approximation and Interpolation: Practical Applications, Challenges and Solutions

Minisymposium Organizers: Athanasios Antoulas, Ion Victor Gosea, Charles Poussot-Vassal

Rational approximation is a fairly established field, which has gone through critical transformations in recent years with the appearance of various reliable and effective algorithms that make computing rational approximations as easy and fast as it has ever been. Many approximation problems from a variety of applications (fluid dynamics, electromagnetics, electrical, mechanical, or aerospace engineering) can now be solved in a matter of seconds (even on personal computing devices), yielding a high level of accuracy. The list of algorithms and developments of rational approximation methods targeted in this MS includes, but is not limited to, methods based on processing Loewner matrices, such as the Loewner framework by [Mayo/Antoulas '07] and the Antoulas-Anderson algorithm (AAA) by [Nakatsukasa/Sete/Trefethen '18]. Such methods, as in the former case, extract relevant information from the measurements by means of an SVD (singular value decomposition) compression, yielding potent approximants through (approximate) interpolation. The latter one blends explicit interpolation with least-squares (LS) fitting in an iterative and adaptive fashion. Another class of methods targeted here solely relies on LS fitting, avoiding interpolation altogether (here, we mention the vector fitting approach, various methods based on optimization through nonlinear LS approaches, VARPRO, etc., or more recent methods based on deep learning). Special attention will be devoted to multivariate rational approximation, including parametric ROMs construction and tensor approximations. Additionally, non-intrusive approaches that target reduced-order modeling of complex dynamical systems are of particular interest.

List of Talks

A refined nonlinear least-squares method for the rational approximation problem

Authors: Linus Balicki[1], Michael Ackermann[1], Serkan Gugercin[1], Steffen W. R. Werner[1]

[1] *Virginia Tech*

Contact: amike98@vt.edu

The adaptive Anderson-Antoulas (AAA) algorithm is capable of generating highly accurate rational approximations to given data. Though AAA almost always produces an approximation to a given target accuracy, the degree of the resulting rational function may be larger than actually required to meet the accuracy tolerance. In this talk, we introduce the nonlinear least-squares adaptive Anderson-Antoulas (NL-AAA) algorithm, which aims to solve the nonlinear least-squares problem arising in the AAA algorithm, as opposed to the linear approximation solved in AAA. The nonlinear problem is solved efficiently with iteratively reweighed least-squares methods. In addition to better accuracy at lower degrees, solving the nonlinear least-squares problem allows us to guarantee monotonic convergence of the NL-AAA method. Further, we provide an analysis of the gradients of each minimization problem, which gives insight into scenarios where AAA is observed to converge sub optimally. Finally, we test our algorithm on numerical examples including classic function approximation problems and applications to reduced order modeling, a field where attaining high accuracy with minimal degree is required.

Descriptor realizations of multi-parameter systems and nonlinear eigenvalue problems

Authors: Athanasios Antoulas[1], Charles Poussot-Vassal, Victor Gosea

[1] *Rice University*

Contact: aca@rice.edu

Purpose of this presentation is to discuss novel descriptor realizations of linear multiple-parameter systems and their connection to nonlinear eigenvalue problems. The work is based on recent developments of the Loewner framework as exposed in a SIAM Review paper, published in November 2025.

Multivariate Rational Approximation of Scattered Data Using the p-AAA Algorithm

Authors: Linus Balicki[1], Serkan Gugercin[2]

[1] *Novateur Research Solutions*

[2] *Virginia Tech*

Contact: balicki@vt.edu

The parametric adaptive Antoulas–Anderson (p-AAA) algorithm is an effective method for multivariate rational approximation [Carracedo Rodriguez et al., 2023], inspired by the AAA framework for univariate rational approximation [Nakatsukasa et al., 2018]. In its original formulation, p-AAA aims to approximate a function $\mathbf{f} : \mathbb{C}^d \to \mathbb{C}$ via a multivariate rational function $\mathbf{r} \approx \mathbf{f}$ based on the samples

$$\mathbf{D} = \{\mathbf{f}(z^{(1)}, \ldots, z^{(d)}) \mid (z^{(1)}, \ldots, z^{(d)}) \in \mathbf{S}\},$$

with the sampling points given by the Cartesian grid

$$\mathbf{S} = \{z_1^{(1)}, \ldots, z_{n_1}^{(1)}\} \times \cdots \times \{z_1^{(d)}, \ldots, z_{n_d}^{(d)}\} \subset \mathbb{C}^d.$$

The main idea of p-AAA is to iteratively construct $\mathbf{r}$ such that it interpolates $\mathbf{f}$ on a subgrid of $\mathbf{S}$ and minimizes the least-squares error over the remaining data. Despite its effectiveness, the original p-AAA algorithm can not always be applied in practice due to its grid data requirement, which may not always be available or feasible to obtain in practice. In this talk, we discuss how the p-AAA algorithm can be formulated with sampling points given by scattered data sets, i.e., arbitrary subsets of $\mathbb{C}^d$ that do not necessarily follow a grid structure. Towards this goal, we introduce a novel approach for enforcing interpolation conditions on scattered data sets via multivariate rational functions in barycentric form. To incorporate this approach into the p-AAA framework, a constrained linear least-squares problem is integrated into the algorithm. We show that this novel formulation of p-AAA is a strict generalization of the original algorithm, and demonstrate its effectiveness on several challenging examples.

Rational quasi-Hermite approximation for computing acoustic quasiresonances in transmission problems

Authors: Luka Grubisic[1], Luka Marohnić[2]

[1] *University of Zagreb, Faculty of Science, Department of Mathematics*

[2] *Zagreb University of Applied Sciences*

Contact: marohnicluka@gmail.com

We introduce an AAA-type method for rational quasi-Hermite approximation formulated in barycentric form. A stacked Hermite–Löwner matrix is assembled from function values and derivative data

at adaptively selected support nodes, and the barycentric weights are determined through a homogeneous least-squares procedure. This approach eliminates the need for external test points as required in the classical AAA algorithm. To maintain computational efficiency, the barycentric weights are computed using randomized nullspace methods, which accelerate the SVD of the stacked Hermite–Löwner matrices.

To further control computational costs and mitigate the appearance of spurious poles, the support set is adaptively partitioned into pieces. This yields piecewise rational quasi-Hermite approximants where inter-piece smoothness is enforced through penalization rather than exact spline constraints, significantly reducing the size of the underlying SVD problems and suppressing unwanted poles. The refinement strategy estimates the location of the maximal residual by constructing a Hermite interpolant of the residual data over the current support nodes. A complete MATLAB implementation is provided, and numerical experiments demonstrate that the method achieves accuracy and efficiency comparable to continuum AAA while requiring substantially fewer function evaluations when derivative information is available.

---

p-set valued AAA for parametric model order reduction

Authors: Giovanni Conni, Karl Meerbergen[1]

[1] *KU Leuven*

Contact: karl.meerbergen@kuleuven.be

We present an extension of the AAA algorithm, named psvAAA. This MOR method combines the multivariate pAAA and uni-variate set valued AAA method. Many physical systems are described by dynamical systems with physical or geometrical parameters. Often, the system's output is strongly dependent on the Laplace variable or the frequency, and less strong on the parameters. In this talks, we present a combination of (p)AAA and set-valued AAA. The 'set' is an ensemble of frequency dependent functions evaluated in points in the parameter domain. This set is used to choose support points. In the talk, we describe ways to determine the weights and derive reduced models for evaluation in frequency or time domain.

Greedy rational approximation: Analysis and algorithms of sketched resolvents

Authors: Akil Narayan[1], Amy DeCastro, Fengyan Li, Filip Belik, Serkan Gugercin, Yanlai Chen

[1] *University of Utah*

Contact: akil@sci.utah.edu

We discuss analytical estimates for greedy construction of rational approximations, where the underlying function is the linear sketch of an operator resolvent. The canonical example of this setup

is the transfer function of a linear dynamical system. Under a sectorial assumption for the operator, this analysis immediately reveals corresponding algorithms, and provides explicit estimates of transfer function rational approximability. The theoretical constructions correspond to explicit Galerkin projection-based reduced order models of the original system, and additionally provide concrete estimates of Hankel singular value decay. The framework we discuss is a paradigm that forges concrete relationships between rational approximation, Kolmogorov n-widths, and Galerkin projection-based model reduction. We will discuss extensions to parametric systems, where this corresponds to multivariate rational approximation through a greedy approach. Finally, we will identify algorithmic challenges and opportunities, investigating numerical procedures that improve on stability and accuracy of the procedure.

The Loewner Framework Beyond Linear Outputs

Author: Sean Reiter[1]

[1] *Courant Institute of Mathematical Sciences, New York University*

Contact: s.reiter@nyu.edu

In recent years, structured reduced-order modeling has become an essential component in meaningful applications across engineering and the physical sciences whenever mathematical models are unavailable, but input-output data are abundant. For linear time-invariant systems, the Loewner framework provides a non-intrusive methodology for the construction of minimal interpolants from rational transfer function data. Its key ingredients are the Loewner and shifted Loewner matrices, which are constructed solely from transfer function evaluations. Significantly, the rank of the Loewner matrix reveals the underlying order (McMillan degree) of the system that generated the data, enabling a trade-off between the accuracy-of-fit and the complexity of the data-based reduced-order model.

In this work, we generalize the Loewner framework to the class of linear quadratic-output systems

$$E\dot{x}(t) = Ax(t) + Bu(t), \quad y(t) = M(x(t) \otimes x(t)).$$

Dynamical systems with quadratic-output functions arise naturally in applications where one is interested in observing or simulating response quantities computed as the product of time- or frequency-components of the state, such as in vibro-acoustic problems or energy-based modeling. We introduce appropriately defined Loewner and shifted Loewner matrices that are built from samples of the multivariate rational transfer function of the underlying quadratic-output model. These matrices retain the hallmark features of the linear time-invariant Loewner approach; namely, their rank exposes the order of the system used to generate the data, and the reduced-order model computed using the proposed Loewner and shifted Loewner matrices is guaranteed to satisfy certain multivariate rational interpolation conditions. We also discuss practical considerations, such as obtaining quadratic-output transfer function data from frequency-response measurements of an equivalent, multiple-output linear time-invariant system.

Concentrated real-pole uniform-in-time approximation of the matrix exponential

Authors: Shuai Shao[1], Stefan Guettel[1]

[1] *The University of Manchester*

Contact: e77968ss@student.manchester.ac.uk

We propose an asympotically optimal choice of shared concentrated real poles of a family of rational approximants of time-dependent exponential functions $\exp(-tz)$ for $z \geq 0$ and t in a positive time interval T. Our result extends a classical result by J.-E. Andersson [J. Approx. Theory, 32(2):85–95, 1981] on the asymptotic best rational approximation of $\exp(-z)$ with real poles. Numerical experiments demonstrate the near-optimality of our choice for various time ranges and for both small and large approximation degrees.

Tensor-based multivariate rational approximation

Authors: Daan Huybrechs[1], Karl Meerbergen[1], Leonie Van Pottelberghe[1]

[1] *KU Leuven*

Contact: leonie.vanpottelberghe@kuleuven.be

Recent contributions for rational approximation include the p-AAA method (and variations) and the extension of the Loewner framework to multiple dimensions. These contributions are recent and are inspiration for further analysis and algorithmic improvements.

In this talk, we combine two ingredients that have proven to be successful in their respective contexts, i.e., the AAA method for univariate rational approximation, and tensor decompositions for the representation of low rank discrete multidimensional data. We present a tensor-based framework for multivariate rational approximation from function samples given on tensor-product grids. The approach combines a low-rank tensor decomposition with the set-valued AAA algorithm. We compare the proposed method against the current state-of-the-art p-AAA by numerical experiments. In particular, we examine both robustness and computational cost on challenging function classes, including nonsmooth and highly oscillatory functions. We also investigate performance in the presence of noise in the sampled data.

MS 26: Advances and Challenges in Eigensolvers

Minisymposium Organizers: Daniel Bielich, Francoise Tisseur

Eigenvalue problems are fundamental to large-scale modelling and simulation across numerous engineering and scientific domains, including acoustics, structural analysis, electromagnetics, and fluid-structure interaction. This minisymposium brings together researchers and practitioners from academia, national laboratories, and industry to showcase recent advances in eigensolver algorithms and to highlight emerging needs posed by next-generation applications. The talks will cover new algorithms, robust and scalable software, and methodological innovations, as well as open challenges.

List of Talks

Integration of the GPLHR Method Within LS-DYNA

Author: Daniel Bielich[1]

[1] *Synopsys*

Contact: dbielich@synopsys.com

LS-DYNA is a multiphysics simulation software package. It targets a wide range of industrial applications, such as modal analysis problems. These are solved with a variety of homegrown eigensolvers, tailored over decades to industrial models. In this talk, we will focus our attention on the quadratic eigenvalue problem underlying the rotational dynamics' framework of the Jeffcott Rotor model. LS-DYNA currently relies on the implementation of IRAM within ARPACK to target this problem. As an alternative, we have developed an in-house extension of the Locally Optimal

Block Preconditioned Conjugate Gradient (LOBPCG) method to the complex case, known as the Generalized Preconditioned Locally Harmonic Residual (GPLHR) method. We present our studies on the numerical behavior and performance of GPLHR, including strong-scaling results (with comparisons to LOBPCG where applicable), sensitivity to tolerances and subspace dimension size, and convergence bottlenecks that trade off preconditioner cost against orthogonalization. Finally, we verify the correctness of the implemented damping contributions by isolating various damping terms in the equation of motion, demonstrating consistent eigenvalue behavior across these cases.

Inexact subspace projection methods for tensor eigenvalue problems

Authors: Alec Dektor[1], Chao Yang, Erika Ye, Peter DelMastro, Roel Van Beeumen

[1] *Lawrence Berkeley National Lab*

Contact: adektor@lbl.gov

I will present two projection methods for solving high-dimensional tensor eigenvalue problems with low-rank structure: an inexact Lanczos method and an inexact polynomial-filtered subspace iteration. The inexactness arises from low-rank compression, which enables efficient representation of high-dimensional vectors in low-rank tensor formats. A central challenge is that standard operations, such as matrix–vector products, increase tensor rank, necessitating rank truncation and thereby introducing approximation errors. Our numerical results show that subspace iteration is significantly more robust to truncation errors than the Lanczos method. Comparisons with the density matrix renormalization group (DMRG) method further demonstrate that subspace iteration can converge in regimes where DMRG stagnates. Overall, these results highlight inexact subspace iteration as a robust and effective approach for computing multiple eigenpairs of rank-structured tensor operators. A convergence analysis of inexact subspace iteration will also be presented.

Adaptive filtered subspace iteration for self adjoint eigenvalue problems on moving domains

Author: Luka Grubisic[1]

[1] *University of Zagreb, Faculty of Science, Department of Mathematics*

Contact: luka.grubisic@math.hr

Filtered subspace iterations can be used to approximate a finite cluster of eigenvalues of a lower semi-bounded selfadjoint operator in a Hilbert space. Prototype examples of such operators are Schrödinger operators with short-range potentials. A rational function (filter) of the operator is constructed such that the eigenspace of interest (eigenvalues below the infimum of the essential spectrum) is its dominant eigenspace, and a subspace iteration procedure is used to approximate this eigenspace. To approximate an operator in an unbounded domain we use a sequence of finitely

truncated domains whose union is the whole space. We present an adaptive multispace algorithm based on *a posteriori* error estimation. Numerical experiments with spectral and finite element approximation methods confirm the theoretical results. We also discuss an application of the results for other moving domain eigenvalue problems.

The Restarted Block Two-Level Orthogonal Arnoldi Algorithm

Author: Michael Jones
Contact: michael.jones.qwe@gmail.com

The two-level orthogonal Arnoldi algorithm, abbreviated as TOAR, proposed by Lu, Su and Bai, is a Krylov method for the solution of large sparse quadratic eigenvalue problems (QEPs). Traditionally, such eigenproblems are first linearised with an appropriate companion form, then fed into the standard Arnoldi algorithm. This approach has the advantage of being simple, but suffers from large memory requirements when storing the basis. In contrast, TOAR works compactly, requiring (to highest order) half the memory of Arnoldi. Similar algorithms existed prior to TOAR, for example the SOAR (second-order Arnoldi) algorithm, which had the same major benefit as TOAR but was prone to numerical instability.

Interest has been shown in block Arnoldi methods for some time because of their ability to exploit BLAS 3 subroutines. Standard block Arnoldi is widely used and some block SOAR algorithms have shown up in the literature, but not block TOAR. In my work I have extended TOAR to block form and have shown that it has similar numerical stability to non-block TOAR.

Stewart's Krylov-Schur algorithm for restarting Arnoldi is an improvement on the implicitly restarted Arnoldi method (IRAM) in terms of numerical stability. Campos and Roman provide an extension of Krylov-Schur to the TOAR algorithm, which I have further extended to block TOAR.

In this talk, I will briefly explain the TOAR algorithm, then cover my work extending TOAR to block form and talk about a numerical instability issue with the restart method.

From eigenvector nonlinearities to eigenvalue nonlinearities

Authors: Vilhelm P. Lithell[1], Elias Jarlebring[1]

[1] *KTH Royal Institute of Technology*

Contact: vipl@kth.se

We are interested in an eigenvector-nonlinear eigenvalue problem (NEPv), that is, a problem of the form $A(x)x = \lambda x$, where $A : \mathbb{R}^n \to \mathbb{R}^{n\times n}$ is symmetric, and the eigenvector has a prescribed norm, for instance $\|x\| = 1$. In this sense, this class of problems generalize the linear eigenvalue problem. Motivated by applications such as the Gross-Pitaevskii equation (GPE) [E. Jarlebring

and P. Henning. SIAM Review, 67:256-317, 2025], we further specialize this problem by considering a specific structure of $A(x)$. More precisely, we consider the case where the nonlinearity appears as a sum of products of scalar functions of the eigenvector, and rank-one matrices. We present a method based on transforming this class of problems into equivalent problems with eigenvalue nonlinearities (NEP), i.e., problems of the form $M(\lambda)x = 0$. In our application, $M : \mathbb{R} \to \mathbb{R}^{n\times n}$ is an algebraic function of the eigenvalue. This transformation enables us to use efficient methods for NEPs as a means of obtaining solutions to the NEPv. In particular, these methods can efficiently compute several eigenvalues, something that is rare among methods for the NEPv. We show how the transformation is constructed theoretically, and how it can be handled in practice. A numerical illustration on a large-scale problem related to a modification of the GPE shows the effectiveness of our approach in computing several eigenvalues of the original problem. This presentation is based on the preprint [E. Jarlebring and V. P. Lithell. arXiv:2506.16182, 2025].

From SCF to LOBPCG: Accelerated Solvers for Nonlinear Eigenvector Problems

Authors: Agnieszka Miedlar[1], Ning Wan

[1] *Virginia Tech*

Contact: amiedlar@vt.edu

This talk surveys modern iterative algorithms for solving large-scale nonlinear eigenvector problems, with a focus on both foundational methods and recent advances in nonlinear acceleration techniques. We examine SCF, FEAST, BPSD, and the LOBPCG eigenvector problems solvers, highlighting their theoretical foundations, practical performance, and applications. Through comparisons and illustrative examples, the talk provides a unified perspective and guidance for selecting and tuning solvers in practice.

The accuracy of the QZ algorithm and some tricks to improve it

Authors: Thijs Steel[1], Jens Verherstraeten[1], Julien Langou[2], Raf Vandebril[1]

[1] *KU Leuven*

[2] *UC Denver*

Contact: thijs.steel@kuleuven.be

For small and/or dense generalized eigenvalue problems, the QZ method remains the method of choice because of its robustness. In this talk, we will discuss how that robustness may occasionally fail, especially if the matrices are badly scaled or singular. We will discuss how new deflation criteria and some changes in the preprocessing can improve the robustness.

Fast Algorithms for Optimal Damping in Mechanical Systems

Authors: Francoise Tisseur[1], Qingna Li[2]

[1] *The University of Manchester*

[2] *Beijing Institute of Technology*

Contact: francoise.tisseur@manchester.ac.uk

Optimal damping aims at determining a vector of damping coefficients ν that maximizes the decay rate of a mechanical system's response. This problem can be formulated as the minimization of the trace of the solution of a Lyapunov equation whose coefficient matrix depends on ν. For physical relevance, the damping coefficients must be nonnegative and the resulting system must be asymptotically stable. We derive explicit expressions for the gradient and Hessian of the objective function and show that the Karush–Kuhn–Tucker conditions are equivalent to the vanishing of a nonlinear residual function at an optimal solution. To compute such a solution, we propose a Barzilai–Borwein residual minimization algorithm (BBRMA), which is simple and efficient but not globally convergent, and a spectral projected gradient (SPG) method, which is globally convergent. By exploiting the structure of the problem, we show how to efficiently compute the objective function and its gradient, with eigenvalue decompositions constituting the dominant cost in terms of execution time. Numerical experiments show that both methods require fewer eigenvalue decompositions than the fast optimal damping algorithm (FODA), and that, although SPG may incur additional decompositions due to line search, it often converges faster than BBRMA, leading to a lower overall computational effort.

MS 27: Linear Algebra Education

Minisymposium Organizers: Jephian C.-H. Lin, Anthony Cronin, Fernando de Terán Vergara

This mini-symposium will highlight innovative and creative ways to engage students in learning linear algebra at all levels, as well as current challenges and opportunities in teaching linear algebra. The ILAS community has an extraordinary wealth and diversity of expertise and experience in education, and our mini-symposium presents an opportunity to share ideas and practices in an international context.

List of Talks

Engaging Undergraduates in Computational Linear Algebra Through Data-Driven Projects

Author: Sima Ahsani[1]

[1] *Washington and Lee University*

Contact: ahsani.sima@gmail.com

This talk will discuss the integration of undergraduate research and project-based learning in both abstract and computational linear algebra. Specifically, it will highlight a research training course focused on the mathematical foundations of machine learning and the analysis of human brain connectivity.

Throughout a summer research program and a Python-based numerical linear algebra course, students are trained to apply computational linear algebra, matrices, and graph theory concepts to real-world datasets. The central project involves investigating structural brain networks derived from 90×90 connectivity matrices obtained from MRI data. Furthermore, students analyze the spectra of adjacency matrices and use computational linear algebra tools to identify structural patterns and hubs within human brain networks.

We will discuss that by engaging students in data-driven projects, students develop proficiency in numerical methods in linear algebra while learning the concept of eigenvalues, eigenvectors, and graph representations. Also, this presentation will reflect on student learning outcomes and the pedagogical benefits of using data to fill the gap between classroom theory and interdisciplinary research.

Linear algebra applications in students' post linear algebra course work

Authors: Christine Andrews-Larson[1], Oluwatosin Deborah Akande[1], Shallamar Goodwin-Richards[1]

[1] *Florida State University*

Contact: cjlarson@fsu.edu

Linear algebra enrollments are increasingly comprised of students from applied STEM majors. Little research has systematically examined the ways in which the curricular resources commonly used in first linear algebra courses relate to the linear algebra applications encountered by students from applied fields in their subsequent coursework and career fields. In this talk, we will summarize findings from interviews with instructors of STEM classes that list a first course in linear algebra as a pre-requisite course. Major areas of focus include computer engineering, electrical engineering, industrial engineering, economics, statistics, computational science, and biomathematics. Interviews focus applied STEM fields instructors' perceptions of how linear algebra is applied in their fields, topics they need to review or teach to cover their own course content, which concepts are most important more broadly in their field, and the kinds of notation and computational tools needed

and used in their disciplines. We also ask instructors to reflect on example applications problems taken from linear algebra textbooks and reflect on those from the lens of their own discipline and students' learning in that discipline. We will provide a preliminary summary of findings from our interviews, relate these to current research on the teaching and learning of linear algebra, and discuss how our findings can inform curriculum development and enactment in the context of linear algebra.

Cards on the Table: Playing with Linear Algebra Concepts

Authors: Martijn Boussé[1], Monica Salvioli[1], Philippe Dreesen[1], Max Sondag[1], Ralf Peeters[1]

[1] *Maastricht University*

Contact: m.bousse@maastrichtuniversity.nl

Linear Algebra is a core mathematics course in numerous undergraduate programs, yet students often struggle with its abstract nature and with recognizing key connections between topics. This lack of conceptual coherence motivates the need for innovative teaching approaches that make linear algebra more accessible, engaging, and meaningful.

We developed a Linear Algebra card game grounded in ludodidactical principles and aimed at fostering conceptual understanding through collaborative learning. The game is not intended to trivialize mathematical content, but to elicit meaningful interaction and explanation among students. It is played in groups of 2 to 6 students and invites players to strategically place and connect concept cards representing key topics in linear algebra (such as systems of linear equations, vector spaces, eigenvalues) by explaining the mathematical relationships between them. Players earn points for correct and meaningful connections, shifting the emphasis from computation to explanation and reasoning.

The game was implemented in a variety of educational settings and evaluated using end-of-course student surveys focusing on engagement and learning outcomes. Preliminary results indicate increased student engagement, improved understanding of relationships between linear algebra concepts, and positive perceptions of learning through peer interaction. In addition, multiple playtesting sessions with colleagues in mathematics education were organized to iteratively refine both the game mechanics and its pedagogical focus.

Our findings suggest that well-designed ludodidactical tools can play a valuable role in supporting conceptual learning in abstract mathematics. The Linear Algebra Card Game is adaptable across different course levels, and all materials are shared as open educational resources.

Visualizing the Spectral Theorem for Symmetric Matrices in a Dynamic Geometry Environment

Author: Günhan Caglayan[1]

[1] *New Jersey City University, Mathematics Department*

Contact: gcaglayan@njcu.edu

Symmetric (real) and Hermitian (complex) matrices occupy a central role in linear algebra due to their well-known spectral properties, including real eigenvalues and mutually orthogonal eigenvectors. While these results are theoretically elegant, students often struggle to develop geometric intuition for eigenvalues, eigenvectors, and matrix transformations, especially in higher dimensions.

This talk presents a sequence of interactive visualizations of symmetric matrices using a dynamic geometry-assisted learning environment, illustrating how eigenvalues and eigenvectors behave in both two and three dimensions. Four representative cases are examined: a 2×2 invertible symmetric matrix with distinct eigenvalues, a 2×2 singular symmetric matrix with a zero eigenvalue, a 3×3 invertible symmetric matrix with three distinct eigenvalues, and a 3×3 singular symmetric matrix with repeated eigenvalues.

Through these examples, we visually demonstrate key theoretical properties of symmetric matrices, including the parallelism between eigenvectors and their matrix images, the scalar-multiplication effect of eigenvalues, and the orthogonality of eigenvectors corresponding to distinct eigenvalues. Special attention is given to the geometric interpretation of zero and repeated eigenvalues, highlighting how singular matrices collapse eigenvectors to the zero vector while preserving orthogonality. These visualizations directly illustrate the Spectral Theorem for symmetric matrices by making eigenvalues, eigenvectors, and orthogonal diagonalization visible through dynamic geometry software. By linking algebraic computations with concrete geometric transformations, the approach provides an intuitive and conceptually coherent framework for understanding the theorem, making it particularly effective for teaching and learning linear algebra.

A Free Online Linear Algebra Textbook with Explorations that may Help your Students

Authors: Anna Davis[1], Paul Zachlin[2]

[1] *The Ohio State University*

[2] *Lakeland Community College*

Contacts: davis.8189@osu.edu, pzachlin@lakelandcc.edu

In January 2023, the authors published the second edition of Linear Algebra: An Interactive Introduction, which the speaker shared at the ILAS Education Session at JMM in Boston.

Thanks to a $ 2.1 million grant from the United States Department of Education entitled "Fortifying Open Education: Scaling Ximera for Enduring Impact", the authors have been able to enhance the online textbook in numerous ways. The book is now on its third edition, and the authors hope to release a fourth edition this fall.

This talk will showcase the Explorations which utilize the capabilities of Ximera. Some of the Octave exercises will also be demonstrated. Instructors who use other textbooks may benefit from using one or two of these activities in their classroom.

Linear Algebra: An Interactive Introduction can be accessed at https://sites.google.com/view/lin-alg-interactive-intro/home

---

Promoting Linear Algebraic Reasoning Among Students: Affordances and Challenges

Author: Guershon Harel[1]

[1] *University of California, San Diego*

Contact: gharel@ucsd.edu

As students transition from the mathematics they learn in school years, including their first-year calculus courses, to the first course in linear algebra, they experience discontinuities in their perspective of what mathematics is. Their propensity to continue applying the same habits of learning in the face of this change leads to failure and frustration. The failure manifests itself in the quality of understanding basic concepts as well as in the lack of linear-algebraic reasoning. Instructional treatments applied in my teaching experiments to foster students' ability to reason linear algebraically resulted in mixed success—some of the treatments were successful, others less so. The latter are accounted for by the structural complexity of the subject matter and students' background knowledge. The pedagogical approaches offered in this talk are oriented within a particular theoretical framework for the learning and teaching of mathematics, called DNR-based instruction in mathematics. Reflections and broader implications are addressed through the lenses of the DNR framework.

You sunk my Battleship! Exploring Matrix Multiplication with a Linear Algebra Video Game

Authors: Matthew Mauntel[1], Bitna Choi[2], Michelle Zandieh[3]

[1] *Tarleton State University*

[2] *University of New Hampshire*

[3] *Arizona State University*

Contact: mmauntel@tarleton.edu

In this talk, we will discuss a game within the video game Vector Unknown:Echelon Seas (VUES) that focuses on matrix multiplication. VUES was designed by a group of students at Arizona State University in collaboration with mathematics educators and computer scientists. The 3D video game allows students to select one of four matrices (represented by cannonballs) and one of four vectors (represented by a cannon) and fire the cannonball at a ship (represented with a target vector) to practice and interpret matrix multiplication that could be leveraged in a first-year linear algebra course. In addition to a symbolic matrix equation, players are also presented with a geometric representation of the matrix as a lattice generated by the column vectors of the chosen matrix.

We conducted several paired teaching experiments with three students from a northeastern university who had already taken linear algebra. During the paired teaching experiments we asked the students to play the game and explain their thinking as they progressed through the level and asked follow-up questions when clarification was needed. We analyzed the results qualitatively and developed a list of student strategies which we highlight during the talk. We divided these strategies into two categories depending on the information from the game the students leveraged for their reasoning: numeric and geometric. One example a student's numeric strategy was to selective compute an part of the matrix multiplication to eliminate potential combinations of matrices and vectors by focusing on a coordinate in the target vector. Another strategy involved students leveraging the geometric supports provided to check the result of the multiplication.

In this talk, we discuss the student strategies we observed and they contributed to students not only learning about matrix multiplication computationally, but also deepening their understanding of the underlying geometric at play.

Linear Algebra Education from an APOS perspective

Author: Asuman Oktaç[1]

[1] *Cinvestav*

Contact: oktac@cinvestav.mx

In the first part of the talk, I will briefly explain the APOS (Action—Process—Object—Schema) theory and how it is applied to the teaching of mathematical concepts as well as to research on learning. In this approach, the mental construction of a mathematical concept is modeled by means of structures that give the theory its name, and mental mechanisms that allow the passage to new structures. The framework helps interpret sources of student difficulties and intuitive ideas in terms of the elements of the theory. It also offers pedagogical suggestions to overcome the challenges.

In our research group, special importance is given to the design of mathematical tasks that promote progress in the learning spiral described by APOS theory. These activities also have the purpose of identifying and differentiating between different mental structures. I will give examples of such

questions involving the concept of linear transformation together with empirical data from students' productions while engaging with the problems. Characteristics of Action, Process and Object conceptions will be discussed both from the point of view of design principles as well as actual student work.

A phenomenon that often happens in linear algebra classes, which can also be observed in textbooks, is that students are expected to deal with Processes and Objects too soon in the instructional development, while they are still constructing Actions. With suitable pedagogical strategies this situation can be dealt with so that students have the opportunity to develop the necessary conceptions for constructing linear algebra concepts.

Symplectic Linear Algebra in Honors Linear Algebra - A Proposal

Author: Matthew Park[1]

[1] *Virginia Tech*

Contact: mfpark@vt.edu

Symplectic linear algebra has had profound connections in many areas of math and physics. Despite this, the subject remains outside the scope of standard linear algebra education. In this talk, I propose motivating the subject via the study of complex coordinate space as the canonical symplectic vector space and argue its approachability for an undergraduate audience. I argue this by highlighting notable theorems and sketching out a hypothetical curriculum.

Immersive Exercises for Linear Algebra

Author: Conrad Plaut[1]

[1] *University of Tennessee–Mathematics*

Contact: cplaut@utk.edu

Modern technology is most often employed in two ways in math education: to enhance explanations (e.g. animated videos or widgets), or to find solutions and calculate answers. The fact that students sometimes do poorly even when taught by excellent teachers shows the limitations of the first strategy. The second strategy, which has been made incredibly easy by generative AI, is increasingly used by students, to the detriment of their learning (and grades). Learning mathematics requires both immersion in problem-solving, and practice. I will describe immersive exercises (IEs) that I have designed for sophomore-level linear algebra, in which students make problem-solving decisions in real time while the software handles routine calculations and provides immediate feedback. Unencumbered by low-level calculations, students may solve problems rapidly, and practice far more (and more complex) problems. External "black box" computational software is not needed. Students quickly realize that entering individual problem-solving steps in real time makes using AI

solutions less efficient than simply learning to solve the problems and practicing. I have designed 75 IEs that have been coded by software engineers from The Math Atlas, which will be licensed to students to be used with the book *Interactive Linear Algebra* published this year by Taylor and Francis.

Balancing Theory and Application in Numerical Linear Algebra with Modern Computational Tools

Author: Sepideh Stewart[1]

[1] *University of Oklahoma*

Contact: sepidehstewart@ou.edu

This talk describes the design and features of a newly developed undergraduate course in numerical linear algebra and an accompanying instructor-led research study that examines how students engage with theoretical structure and computational behavior in AI-supported environments.

This course offered an introduction to numerical linear algebra, focusing on its theoretical foundations, practical applications, and computational techniques. Topics included matrix factorization (LU, QR, Cholesky, SVD), eigenvalues and eigenvectors, iterative methods, and least squares. Computational explorations in MATLAB were central to the course, allowing students to learn theories through computation and real-world applications, bridging the numerical and abstract linear algebra.

The instructor conducted educational research by collecting students' work, regular reflections, and an exit survey, along with the instructor's daily teaching reflections, to document classroom interactions and instructional decisions. The goal is to examine how students interpret numerical output, reason about algorithmic reliability, and incorporate AI-assisted tools while remaining grounded in theory. Some preliminary data and analysis will be shared in this talk.

The course was designed and created in collaboration with Mike Michailidis and Jon Loftin from MathWorks.

Finding the Right Basis (or Bases)

Author: Jeffrey Stuart[1]

[1] *Pacific Lutheran University*

Contact: jeffrey.stuart@plu.edu

Traditionally, half or more of a first course in linear algebra has been devoted to the theory and machinery required to solve linear systems and to find bases for the four fundamental subspaces

associated with a matrix. At the heart of this is the Gauss-Jordan algorithm, used to construct the reduced row echelon form (RREF) of a matrix. Far too many students complete their only linear algebra by desperately clutching the manual computation of the RREF as if it were a life preserver in a frightening sea of theory, an approach fostered by many texts.

Increasingly, students arrive in linear algebra less well-prepared to handle definitions and the reasoning of proofs. Simultaneously client disciplines are increasingly using concepts rarely or only lightly covered in a first course: the spectral theorem, the QR algorithm, orthogonal projections, the SVD (or its applied cousin, PCA). Data science and machine learning, heavy users of these techniques, did not even exist as programs of study ten years ago.

Over the last decade, talks in linear algebra education sessions have focused on what needs to be in a "new" version of linear algebra that better addresses the changing needs of our client disciplines. Frank Uhlig has been a notable proponent of making significant changes. Others have suggested ways to cleverly fold advanced topics into a traditional course. How do we reconcile the weaknesses of incoming students and the deepening topical needs of our clients?

My original title for this talk was **Matrices Just Want to Be Diagonal**, a title which emphasizes algorithms. Then I realized that those algorithms are really consequences of a different talk, **Finding The Right Basis (or Bases)**. Naturally, I will discuss what constitutes a "right" basis, and what content changes to an introductory course might enable us to develop such bases and the resulting algorithms.

Blind Singal Separation as a model to introduce Linear Transformatios

Authors: Rita Vazquez[1], Maria Trigueros[2]

[1] *Univ Autonoma Ciudad de México*

[2] *Benemérita Universidad Autónoma de Puebla, México*

Contact: mtriguerosg@gmail.com

The use of real or realistic problems to introduce university students to linear algebra concepts through modeling has proven effective in stimulating student learning. We present an experience based on mathematical modeling and APOS theory (Action, Process, Object, and Schema) to introduce the concepts of matrix transformation and inverse matrix transformation in the context of an introductory Linear Algebra course for engineering students.

The instructional sequence is inspired by Blind Source Separation (BSS), a complex engineering problem concerned with recovering original sources from observed mixtures. Building on a prior in-depth study of BSS models, we identify a linear, noise-free formulation in which observations are modeled as linear combinations of sources. From this formulation, we carried out a transposition of the structural elements of the problem to an introductory linear algebra setting, interpreting signals as vectors and mixing processes as linear transformations.

Based on this model, the inverse transformation emerges as a natural response to the associated inverse problem. Rather than introducing the inverse matrix as a formal object or technique, the proposed genetic decomposition describes how the need for an inverse transformation arises from interpreting linear transformations as models of mixing processes. The construction involves the coordination of various mathematical processes, including the matrix–vector product, systems of linear equations, and the interpretation of linear transformations as functions between vector spaces. Invertibility is conceptualized in terms of information loss and the impossibility of reversing certain transformations.

This work contributes to the literature through the design of a modeling-based instructional situation that foregrounds the emergence and construction of inverse transformations in an introductory linear algebra context.

The Inquiry-Oriented Linear Algebra Project

Authors: Megan Wawro[1], David Plaxco[2], Christine Andrews-Larson[3], Michelle Zandieh[4]

[1] *Virginia Tech*

[2] *University of Georgia*

[3] *Florida State University*

[4] *Arizona State University*

Contact: mwawro@vt.edu

The Inquiry-Oriented Linear Algebra (IOLA) project is a research-based, student-centered approach to the teaching and learning of introductory linear algebra. The IOLA curricular materials build from a set of experientially real tasks that allow for active student engagement in the guided reinvention of key mathematical ideas through student and instructor inquiry. The online instructional support materials include various resources such as rationales for task design, implementation suggestions, and examples of student work. In this talk, we will highlight the IOLA least squares task sequence, which builds from a travel metaphor towards students' creation of the normal equation and understanding of why it is relevant in solving the least squares problem.

A few issues in teaching linear algebra

Author: Fuzhen Zhang[1]

[1] *Nova Southeastern University*

Contact: zhang@nova.edu

Linear algebra is fundamental to mathematics and has wide-ranging applications in science, engineering, and data analysis. As a core mathematical subject, it is taught to students at both the

high school and college levels, not only in mathematics but also across the sciences. Is linear algebra easy? It may not be as easy as it sounds. Is teaching linear algebra easy? It may not be as easy as one thinks. Seemingly simple questions often reveal subtle conceptual issues: for example, what is the degree of the zero polynomial, and why? If every subset of a vector space spans a subspace, what does the empty set span? Are the eigenvalues of a matrix continuous functions of its entries?

This talk addresses several issues in linear algebra arising in both teaching and research. Such issues are frequently overlooked or treated carelessly in textbooks and classroom instruction.

MS 28: Symplectic Linear Algebra and Applications

Minisymposium Organizers: Tanvi Jain, Hemant Mishra

Symplectic linear algebraic techniques have profound applications in various areas of physical and mathematical sciences including Hamiltonian dynamics, quantum mechanics, optimization, and symplectic geometry. It is useful in the study of symplectic manifolds, which are even dimensional linear spaces equipped with non-degenerate, skew-symmetric bilinear forms. The structure of symplectic matrices help formalize the concepts of volume and area preserving linear transformation, analogous to the concept of orthogonal matrices preserving length and angle. Symplectic matrices play an essential role in Hamiltonian dynamics, and for model order reduction of Hamiltonian systems. The established importance of symplectic matrices has also led to inquiry into the symplectic Stiefel manifold and optimization problems over this manifold. Symplectic numerical linear algebra and integration techniques have found applications in Hamiltonian optics. There has been much interest in the perturbation theory and Hamiltonian matrices and matrix decompositions involving symplectic matrices. A result of central focus in the standard symplectic space is Williamson's theorem that gives rise to the notion of symplectic eigenvalues. Williamson's theorem plays a key role in developing a comprehensive mathematical formalism of bosonic Gaussian quantum states, making this class of quantum states more accessible to quantum information theorists. This symposium will focus on recent advances and novel methods in symplectic linear algebra, and applications of symplectic techniques in other areas of physics and mathematics. The aim is to unite researchers across various career stages who have employed symplectic methods in their research and to provide a conducive environment for the exchange of ideas.

List of Talks

The ϕ-Reversibility Problem for the Real Symplectic Group

Authors: Nicole Joy Datu[1], Ralph John de la Cruz[1]

[1] *University of the Philippines Diliman*

Contact: njdatu@math.upd.edu.ph

Let G be a matrix group over a field $\mathbb{F}$ and $\phi : M_n(\mathbb{F}) \to M_n(\mathbb{F})$ be a map such that $\phi(A) \in G$ for all $A \in G$. An element $A \in G$ is said to be *$\phi-$reversible* if there exists $P \in G$ such that $PAP^{-1} = \phi(A)$. If P can be chosen to be an involution (i.e., $P^2 = I$), then A is said to be *strongly $\phi-$reversible*. The strongly $\phi-$reversible elements of the complex symplectic group

$$\mathrm{Sp}(2n, \mathbb{C}) := \left\{ A \in GL_{2n}(\mathbb{C}) \,\middle|\, A^T \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix} A = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix} \right\}$$

for the map $\phi : A \mapsto -A$ have been completely classified. In this talk, we consider the real case. We show that several results from the complex setting no longer hold over $\mathbb{R}$. In particular, for each n, we construct a real symplectic matrix that is strongly $\phi-$reversible as an element of $\mathrm{Sp}(4n, \mathbb{C})$, but not as an element of $\mathrm{Sp}(4n, \mathbb{R})$. Using a suitable canonical form for real symplectic matrices, we classify the strongly $\phi-$reversible elements of the real symplectic group $\mathrm{Sp}(4, \mathbb{R})$.

Structure-preserving Krylov Subspace Approximations for the Matrix Exponential of Hamiltonian Matrices

Authors: Heike Faßbender[1], Michel Senn[1], Peter Benner[2]

[1] *TU Braunschweig, Institute for Numerical Analysis*

[2] *Max Planck Institute for Dynamics of Complex Technical Systems*

Contact: h.fassbender@tu-braunschweig.de

It is well known that the matrix exponential $\exp(H)$ is symplectic whenever $H \in \mathbb{R}^{2n\times 2n}$ is a Hamiltonian matrix. A matrix H is called Hamiltonian if it satisfies $HJ = (HJ)^T$, while a matrix S is called symplectic (or J-orthogonal) if $S^T JS = J$. Here, $J \in \mathbb{R}^{2n\times 2n}$ denotes $J = \left[\begin{smallmatrix} 0 & I_n \\ -I_n & 0 \end{smallmatrix}\right]$.

In this talk, we study structure-preserving Krylov subspace methods for approximating the matrix–vector products $f(H)b$, where H is a large Hamiltonian matrix and f denotes either the matrix exponential or the related φ-function ($\varphi(z) = \frac{e^z-1}{z}$). Such computations are central to exponential integrators for Hamiltonian systems. Assume that a suitable projection $\Pi = VW^T \in \mathbb{R}^{2n\times 2n}$ is given, where $V, W \in \mathbb{R}^{2n\times 2m}$, $m \le n$, have full column rank and $W^T V = I_{2m}$. Then $f(H)b$ for $b \in \mathbb{R}^{2n}$ can be approximated by $f(H)b \approx V f(\widetilde{H}) W^T b$ with $\widetilde{H} = W^T HV \in \mathbb{R}^{2m\times 2m}$. When $m \ll n$, evaluating $f(\widetilde{H})b$ is far less computationally expensive than evaluating $f(H)b$. Typically, $V = W$ and the columns of V form an orthonormal basis of the standard Krylov subspace $\mathcal{K}_{2m}(H, b) = \mathrm{span}\{b, Hb, \ldots, H^{2m-1}b\}$ of order $2m$, but then $\widetilde{H}$ is in general not a Hamiltonian matrix. This motivates the use of Krylov bases with J-orthogonal columns that yield Hamiltonian projected matrices and symplectic reduced exponentials. We compare several such structure-preserving Krylov methods on representative Hamiltonian test problems, focusing on accuracy, efficiency, and

structure preservation, and briefly discuss adaptive strategies for selecting the Krylov subspace dimension.

On generalization of Williamson's theorem to real symmetric matrices

Author: Hemant Mishra[1]

[1] *Indian Institute of Technology (ISM) Dhanbad*

Contact: hemantmishra1124@gmail.com

Williamson's theorem states that if A is a $2n \times 2n$ real symmetric positive definite matrix then there exists a $2n \times 2n$ real symplectic matrix M such that $M^T AM = D \oplus D$, where D is an $n \times n$ diagonal matrix with positive diagonal entries known as the symplectic eigenvalues of A. The theorem is known to be generalized to $2n \times 2n$ real symmetric positive semidefinite matrices whose kernels are symplectic subspaces of $\mathbb{R}^{2n}$, in which case, some of the diagonal entries of D are allowed to be zero. In this work, we further generalize Williamson's theorem to $2n \times 2n$ real symmetric matrices by allowing the diagonal elements of D to be any real numbers, and thus extending the notion of symplectic eigenvalues to real symmetric matrices. Also, we provide an explicit description of symplectic eigenvalues, construct symplectic matrices achieving Williamson's theorem type decomposition, and establish perturbation bounds on symplectic eigenvalues for a class of $2n \times 2n$ real symmetric matrices denoted by EigSpSm$(2n)$. The set EigSpSm$(2n)$ contains the set of $2n \times 2n$ real symmetric positive semidefinite matrices whose kernels are symplectic subspaces of $\mathbb{R}^{2n}$. Our perturbation bounds on symplectic eigenvalues for EigSpSm$(2n)$ generalize known perturbation bounds on symplectic eigenvalues of positive definite matrices given by Bhatia and Jain *[J. Math. Phys. 56, 112201 (2015)]*.

On symplectic reduction of a matrix to upper J-Hessenberg form

Author: Ahmed Salam[1]

[1] *Université du Littoral-Côte d'Opale*

Contact: salam@univ-littoral.fr

In the context of computations of eignevalues and eignevectors, structure-preserving of a class of specific structured matrices, the reduction of a matrix to a J-Hessenberg condensed form is needed. Such reduction is based on symplectic similarity transformations. It is a crucial step in the SR-algorithm (which is a QR-like algorithm), structure-preserving, for computing eigenvalues and vectors of such structured matrices. The algorithm JHESS, or its variant JHMSH are the main algorithms used for such reduction.

These algorithms may meet fatal breakdowns, causing brutal stops of the computations or encounter near-breakdowns, which are source of serious numerical instability.

In this talk, we point out the source of these breakdowns and present strategies for curing them.

MS 29: Low-rank Matrix and Tensor Decompositions: Theory, Algorithms and Applications

Minisymposium Organizers: Subhayan Saha, Stefano Sicilia

Low-rank models such as matrix factorizations and tensor decompositions constitute a unifying paradigm that has fueled extensive research across linear algebra, signal processing, machine learning, to name a few. In these models, observed high-dimensional data are assumed to lie close to a low-dimensional linear or multi-linear subspace. For instance, two of the most well known representations of a matrix or tensor large data are matrix factorization and Canonical Polyadic (CP) Decomposition. Depending on the dataset and the features required by the applications, one may impose further constraints on the factors, such as non-linearity, non-negativity as well as consider other models such as coupled matrix-tensor factorizations, Tucker decomposition, etc.. These models provide interesting insights in data compression and pattern recognition, but they also pose theoretical and computational challenges such as: Identifiability: determine when the factors of the decompositions are essentially unique is critical for interpreting the model; Recoverability and robustness: unveil the underlying low-rank structure and understand how perturbations affect stability. Scalabilty: as the given data grows in dimensionality and volume, algorithms must efficiently compute low-rank approximations under the problem constraints while keeping time and memory costs tractable. This minisymposium aims to showcase research that delves into these theoretical and algorithmic frontiers. Contributions will address the geometry of the low-rank structures, identifiability, robustness and scalability of the problems and will offer new results connecting classical low-rank ideas with the latest challenges in the field.

List of Talks

Fast randomized compression of matrix-vector products in tensor-train format and applications to Krylov subspace methods

Author: Alberto Bucci[1]

[1] *University of Edinburgh*

Contact: abucci2@ed.ac.uk

We present a new technique for efficiently compressing matrix–vector products in the tensor-train (TT) format, avoiding the explicit formation of intermediate tensors arising from standard MPO–MPS multiplication. The proposed method performs the compression in a single pass, leading to significant computational and memory savings.

From a theoretical point of view, the resulting approximation is mathematically equivalent to the streaming tensor-train approximation introduced by Kressner, Vandereycken, and Voorhaar. However, our approach admits a more efficient implementation, tailored to the structure of TT operators and vectors, and better suited for large-scale and iterative computations.

Beyond accelerating the fundamental MPO–MPS product, the new method allows us to produce a compressed representation of the operator that can be computed once and reused across multiple applications. This feature is particularly advantageous in iterative algorithms where the same operator is applied repeatedly. We demonstrate how this leads to substantial speedups in Krylov subspace methods, with a special focus on the synergy between the proposed technique and sketched GMRES.

Randomized algorithms for streaming low-rank approximation in tree tensor network format

Authors: Alberto Bucci[1], Gianfranco Verzella[2]

[1] *University of Edinburgh*

[2] *University of Geneva*

Contact: gianfrancover96@gmail.com

In this work, we present the tree tensor network Nyström (TTNN), an algorithm that extends recent research on streamable tensor approximation, such as for Tucker or tensor-train formats, to the more general tree tensor network format, enabling a unified treatment of various existing methods. Our method retains the key features of the generalized Nyström approximation for matrices, i.e. it is randomized, single-pass, streamable, and cost-effective. Additionally, the structure of the sketching

allows for parallel implementation. We provide a deterministic error bound for the algorithm and, in the specific case of Gaussian dimension reduction maps, also a probabilistic one. We also introduce a sequential variant of the algorithm, referred to as sequential tree tensor network Nyström (STTNN), which offers better performance for dense tensors. Furthermore, both algorithms are well-suited for the recompression or rounding of tensors in the tree tensor network format. Numerical experiments highlight the efficiency and effectiveness of the proposed methods.

Symmetric nonnegative trifactorization rank of matrices with a given pattern without a four cycle

Authors: Damjana Kokol Bukovšek[1], Polona Oblak[1], Helena Šmigoc[2]

[1] *University of Ljubljana*

[2] *University College Dublin*

Contact: damjana.kokol.bukovsek@ef.uni-lj.si

We consider a symmetric nonnegative matrix A of order $n \times n$. A factorization of the form $A = BCB^T$, where B is a nonnegative matrix of order $n \times k$ and C is a symmetric nonnegative matrix of order $k \times k$, is called symmetric nonnegative trifactorization (SNT for short) of A. Minimal possible k in such factorization is called the SNT-rank of A.

The zero-nonzero pattern of a matrix can be described by a simple graph that allows loops. The SNT-rank of a graph G is the minimal SNT-rank of all symmetric matrices with pattern determined by G, and it can be characterized combinatorially using set-join covers of G. In the talk we will consider a family of graphs that do not contain four cycles. We will present an algorithm on the graph for computing SNT-rank of such graphs.

Robustness of Minimum-Volume Nonnegative Matrix Factorization under an Expanded Sufficiently Scattered Condition

Authors: Giovanni Barbarino[1], Nicolas Gillis[2], Subhayan Saha[2]

[1] *Universidad de Huelv*

[2] *Universite de Mons*

Contact: subhayan.saha@umons.ac.be

Minimum-volume nonnegative matrix factorization (min-vol NMF) has been used successfully in many applications, such as hyperspectral imaging, chemical kinetics, spectroscopy, topic modeling, and audio source separation. However, its robustness to noise has been a long-standing open problem. In this paper, we prove that min-vol NMF identifies the groundtruth factors in the presence of noise under a condition referred to as the expanded sufficiently scattered condition

which requires the data points to be sufficiently well scattered in the latent simplex generated by the basis vectors.

Manifold-based Algorithms for the Hadamard Decomposition

Authors: Stefano Sicilia[1], Arnaud Vandaele[1], Nicolas Gillis[2], Subhayan Saha[2]

[1] *University of Mons*

[2] *Universite de Mons*

Contact: stefano.sicilia@umons.ac.be

Given a matrix X, and two ranks r_1 and r_2, the Hadamard decomposition (HD) consists in looking for two low-rank matrices, X_1 of rank r_1 and X_2 of rank r_2, both of the same size as X, such that $X \approx X_1 \circ X_2$, where $\circ$ is the Hadamard (element-wise) product. HD is more expressive than standard low-rank approximations, such as the truncated singular value decomposition (TSVD), as it can represent higher-rank matrices with the same number of parameters; this is because the rank of $X_1 \circ X_2$ is generically equal to $r_1 r_2$. In this paper, we first present some new theoretical results for HD. These allow us to develop two new manifold-based gradient descent algorithms for computing HDs: the first one uses the representation $X \approx X_1 \circ X_2$, while the second one uses $X \approx WH^\top$, where W and H have $r_1 r_2$ columns. We also propose new initializations that allow us to obtain better HDs. We compare our algorithms and initialization strategies with the TSVD and with the state of the art. Numerical results show that the new methods are efficient and competitive, both on synthetic and real data.

Chiselling Algorithms for Algebraic Computation of Tensor Block Term Decompositions

Authors: David Thorsteinsson[1], Nick Vannieuwenhoven[1], Tim Seynnaeve[2]

[1] *KU Leuven*

[2] *Polish Academy of Sciences*

Contact: david.thorsteinsson@kuleuven.be

Block term decomposition (BTD) unifies the two most common tensor decompositions: canonical polyadic and Tucker. While BTDs have found a broad range of applications from machine learning to blind source separation, all known algorithms for computing BTDs were historically optimisation-based, and required the desired block sizes to be specified as input. Recently, algebraic BTD algorithms have been found. Cai & Li (2021) were the first to find such an algorithm, with Domanov et al. (2024) later publishing a second, very similar algorithm. While these algorithms work well, very little was known about their theoretical underpinnings.

In this talk, we explain the underlying mechanism of these algorithms by identifying them as special cases of Brooksbank, Kassabov, & Wilson's Lie-theoretic *chiselling* framework. This perspective allows us to introduce a third, previously unexplored algorithm variant; for third-order tensors, these three variants constitute the only possible chiselling methods for BTD. Using this framework, we characterise the BTDs recoverable by chiselling: specifically, those where the sum of blocks constitutes a direct sum. We further show that for an arbitrary tensor, the set of detectable tensor product subspaces naturally forms a lattice.

Finally, we address uniqueness via two definitions. The first is subspace identifiability, which we show holds generically for all chiselled BTDs. The second is a stronger definition of uniqueness discussed by the previous authors. In particular, this *factor matrix uniqueness* requires not only unique subspaces, but also that each subspace should be detected with a unique set of bases. Regarding this, Cai & Li claimed their algorithm yields generic uniqueness but did not provide a proof, while Domanov et al. provided only sufficient conditions. We resolve these gaps by formally proving Cai & Li's claim and deriving necessary and sufficient conditions for generic uniqueness across all three variants.

Reducing swamp behavior for canonical polyadic decomposition

Authors: Charlotte Vermeylen[1], Lieven De Lathauwer[1], Nico Vervliet[1]

[1] *KU Leuven*

Contact: charlotte.vermeylen@kuleuven.be

A novel optimization framework is proposed for solving the low-rank tensor approximation problem using the canonical polyadic decomposition (CPD). This can be a difficult optimization problem for certain tensors, e.g., due to degeneracy, i.e., a tensor that can be approximated arbitrarily closely by an ill-conditioned tensor of lower rank. This is one of the phenomena that are encountered in regions of slow convergence, informally known as swamps.

Numerical experiments with state-of-the-art optimization algorithms indicate that in a swamp, often only a few rank-1 terms are modified while others stagnate. Often, the non-stagnant terms are problematic and form an ill-conditioned decomposition. To address this, we propose to temporarily freeze the stagnant terms. The lower number of terms has several benefits: it simplifies the problem by reducing the number of variables and reduces the cost per iteration significantly. Furthermore, in many cases, the residual tensor can be compressed, and an algebraic (re)initialization can be carried out, even if this was not possible for the original tensor. A refinement step can further improve the accuracy if desired. We provide theoretical insights into why terms can stagnate. More specifically, we prove that terms that are close to the solution are not modified anymore in further optimization steps under certain assumptions.

Extensive numerical experiments show that our framework greatly facilitates escaping from swamps. The resulting algorithm outperforms current state-of-the-art approaches on difficult-to-decompose

tensors, both in accuracy and computation time, and has similar performance on easier problems.

MS 30: Spectral Interlacing, Graph Learning, and Quantum Perspectives on Signed Graphs

Minisymposium Organizers: Amrita Mandal, Ravi Srivastava

The abstract focuses on the synergy among spectral graph theory, network science, and quantum computation, in line with linear algebra and the broader applications of the subject area. The proposed minisymposium aims to foster researchers in advancing interdisciplinary work with new research methodologies, theoretical understanding, practical implementation, and future research itinerary driven by linear algebraic tools. Spectral features and distance-based representations are widely used in graph neural networks and social network analysis. Understanding how spectral properties behave under graph transformations helps improve the stability, reliability, and fairness of graph-based learning methods. Along with classical spectral results, the minisymposium explores connections between graph theory and quantum systems. Graph Laplacians and distance-based matrices arise naturally in quantum walks and quantum network models, where eigenvalues govern the evolution and interaction of quantum states on graphs. These connections highlight the growing importance of graph spectra in quantum computing and quantum communication networks. The next session covers interlacing results for net-Laplacian, distance Laplacian, and normalized distance Laplacian matrices under edge deletion, vertex removal, contraction, and vertex replication. The session also presents recent corrections and improvements to known interlacing results for normalized matrices and characterization of signed graphs that attain the extremal eigenvalue. The session further highlights structural notions such as k-geodesic μ-balanced signed graphs and reverse oriented (r-oriented) signed graphs. This minisymposium also characterizes balancedness and distance-compatibility of different graph products. Overall, this minisymposium connects spectral theory with graph learning, network analysis, and quantum technologies, motivating further research in spectral signed graph theory.

List of Talks

Quantum-Classical Algorithms for Counting Triangles in a Signed Edge Stream

Authors: Steven Kordonowy[1], Bibhas Adhikari[2], Hannes Leipold[2]

[1] *University of California at Santa Cruz, CA, USA*

[2] *Fujitsu Research of America, Inc.*

Contact: badhikari@fujitsu.com

We discuss a quantum-classical streaming algorithm that processes signed edges to efficiently estimate the counts of triangles of diverse signed configurations in the edge stream. The approach introduces a quantum sketch register for processing the signed-edge-stream, together with measurement operators for query-pair calls in the quantum estimator, while a complementary classical estimator accounts for triangles not captured by the quantum procedure. This hybrid design yields a polynomial space advantage over purely classical approaches, extending known results from unsigned edge-stream data to the signed setting.

State transfer in discrete quantum walks: from coins to weighted graphs

Author: Hanmeng Zhan[1]

[1] *Worcester Polytechnic Institute*

Contact: hzhan@wpi.edu

A discrete quantum walk takes place on the arcs of a graph, and evolves according to a coin operator and a shift operator. One important task, given the underlying graph, is to construct quantum walks that start with a state localized at a vertex and get arbitrarily close to a state localized at another vertex. In this talk, I will show how different coin operators translate into different weightings of the underlying graph, and use the spectral properties of these weighted graphs to construct quantum walks that admit the aforementioned transport phenomenon.

MS 31: Sparse Tensor Computations: Algorithms and Applications

Minisymposium Organizers: Tianyi Shi, Navjot Singh

Tensors are multidimensional generalizations of matrices and arise naturally in applications ranging from discretized PDEs and quantum chemistry to machine learning and data science. However, the curse of dimensionality poses a fundamental computational challenge. Storage and computation costs grow exponentially with the number of dimensions, and even basic operations can quickly become infeasible. Promoting sparsity through low rank structure or elementwise zeros is therefore essential to making tensor computations tractable. Exploiting such sparsity enables efficient algorithms for tensor decomposition, completion, and contraction, and allows for solving large-scale problems on modern parallel architectures. This minisymposium brings together researchers working on mathematical and computational aspects of sparse and data sparse tensors to discuss recent algorithmic advances and their applications.

List of Talks

Optimal tensor network structure search

Authors: Zheng Guo[1], Aditya Deshpande[2], Alex Gorodetsky[2], Brian kiedrowski[2], Xinyu Wang[2]

[1] *University of MIchigan*

[2] *University of Michigan*

Contact: goroda@umich.edu

Tensor networks provide a powerful framework for compressing multi-dimensional data. The optimal tensor network structure for a given data tensor depends on both data characteristics and specific optimality criteria, making tensor network structure search a difficult problem. Existing solutions typically rely on sampling and compressing numerous candidate structures; these procedures are computationally expensive and therefore limiting for practical applications. We address this challenge by viewing tensor network structure search as a program synthesis problem and introducing an efficient constraint-based assessment method that avoids costly tensor decomposition. Specifically, we establish a correspondence between transformation programs and network structures. We also design a novel operation named output-directed splits to reduce the search space without hindering expressiveness. We then propose a synthesis algorithm to identify promising network candidates through constraint solving, and avoid tensor decomposition for all but the most promising candidates. Finally, we extend this approach to work both with cross approximation and for arbitrary reshaping and re-orderings of nodes Experimental results show that our approach improves search speed by up to 10× and achieves compression ratios 1.5× to 3× better than state-of-the-art. Notably, our approach scales to larger tensors that are unattainable by prior work. Furthermore, the discovered topologies generalize well to similar data, yielding compression ratios up to 2.4× better than a generic structure while the runtime remains around 110 seconds.

Efficient oversampled Tensor-Train approximations

Authors: Daniel Hayes[1], Jingmei Qiu[1], Tianyi Shi[2]

[1] *University of Delaware*

[2] *Lawrence Berkeley National Laboratory*

Contact: dphayes@udel.edu

Recently, there have been many advances in the area of randomized and sampling-based methods for data approximation. This has led to significant progress towards the efficient treatment of large data in both compression and utilization in computation. In this talk, I will discuss a current work that uses random oversampling on a Tensor Train Cross (TT-Cross) approximation in order to reduce the observed error of a tensor approximation. This work includes two separate formulations utilizing projection techniques to construct tensor cores with results demonstrating a reduction in error. Along with the observed reduction in error, we will show that in practice, the oversampling procedure does not substantially increase computation time compared to a standard TT-Cross construction.

A Tree Tensor Network Integrator for the Chemical Master Equation

Authors: Julian Mangott[1], Martina Prugger[2], Lukas Einkemmer[1]

[1] *Universität Innsbruck*

[2] *Max-Planck-Institut für Plasmaphysik (Garching)*

Contact: julian.mangott@gmail.com

The development of new drugs and therapies increasingly relies on the numerical simulation of the reaction networks inside biological cells. However, the most accurate description of such reaction networks with the chemical master equation (CME) suffers from the curse of dimensionality, meaning that memory and computational cost grow exponentially with the number of dimensions. This renders the simulation of such networks infeasible on currently available computer hardware and it will remain a fundamental limitation in the future.

In this talk, we present a novel numerical solution method for the stochastic description of chemical reaction networks. It is based on the dynamical low-rank approximation with tree tensor networks, which effectively approximates the high-dimensional solution of the CME as a linear combination of low-rank factors.

This low-rank approximation corresponds to a hierarchical separation of the reaction network into smaller partitions. Only reactions that occur between chemical species in different partitions are approximated. The number of low-rank factors, the so-called rank, determines the accuracy of the approximation; typically, 10 to 20 low-rank functions are sufficient for a good solution.

We derive an efficient integration scheme for the low-rank factors that exploits the low-rank tensor structure of the CME. This scheme is based on the projector-splitting integrator for tree tensor networks [Ceruti, Lubich, and Walach, SIAM J. Numer. Anal., vol. 59, 2021]. We apply our method to models from biochemistry and show that the method significantly reduces the memory consumption, while achieving improved computational performance and better runtimes compared to a Monte Carlo method.

Finally, we investigate how the choice of the partitioning affects the required rank. To this end, a heuristic partitioning algorithm is presented that minimizes information entropy and the number of reactions between partitions.

Synchronous and Asynchronous Parallelism Approaches for Generalized Canonical Polyadic Tensor Decomposition with GenTen

Authors: Eric Phipps[1], Jeremy Myers

[1] *Sandia National Laboratories*

Contact: etphipp@sandia.gov

The Canonical Polyadic (CP) tensor decomposition is a well-known method for interpretable analysis of high-dimensional data. Recently, the Generalized CP (GCP) method was introduced by Hong, Kolda and Duersch (2020) to allow for flexible choice of the loss function in the optimization problem defining the CP model, enabling more interpretable decompositions of strongly non-Gaussian data such as count or binary data. Furthermore, Kolda and Hong (2020) introduced a version

of GCP that leverages randomization and stochastic optimization to address scalability to large, sparse data sets. In this work, we take these ideas a step further and consider synchronous and asynchronous algorithms for parallel GCP tensor decomposition through the GenTen software package, exploiting both shared and distributed memory parallelism. We build on shared memory parallel CP decomposition algorithms utilizing Kokkos for portability across CPU and GPU architectures (Phipps and Kolda, 2019) to support the random sampling and stochastic optimization methods required by GCP. We then couple this approach to the well-known medium-grained distributed memory parallelism scheme (Smith and Karypis, 2016) developed for traditional CP decompositions through MPI, providing a synchronous, hybrid MPI+Kokkos, parallel GCP decomposition capability. Finally, we propose an asynchronous distributed parallelism approach building on related techniques for federated learning to achieve even better scalability to large data sets. We study the effectiveness of the proposed synchronous and asynchronous approaches vis-à-vis computational cost and accuracy on synthetic and publicly available real-world datasets of varying sizes, dimensions, and sparsity patterns using several loss functions.

Iterative low-rank time integration of the time-dependent Schrödinger equation

Authors: Federico Vismara, Markus Bachmayr, Matthieu Dolbeault, Polina Sachsenmaier[1], Tianyu Jin

[1] *RWTH Aachen University*

Contact: sachsenmaier@igpm.rwth-aachen.de

Standard numerical methods for solving PDEs typically suffer from the curse of dimensionality: their computational cost scales exponentially with the dimension of the underlying domain, making them impractical even at low resolution. In many cases of interest, however, such limitations can be overcome by appropriately compressed representations of approximate solutions, for example, by low-rank tensor representations.

For time-dependent PDEs, several approaches to low-rank approximation exist that control ranks in different ways. These range from methods that keep ranks fixed, such as dynamical low-rank approximations, which may lead to uncontrolled errors, to methods that approximate standard time-stepping schemes to any desired accuracy but can produce unnecessarily large ranks.

We develop time integration methods in low-rank tensor representations that adaptively adjust the approximation ranks to meet a prescribed accuracy while simultaneously maintaining control over the ranks of the computed approximations and all intermediate quantities. These ranks are a main determining factor in the computational costs of such methods. Our approach combines an iterative time-stepping scheme with soft thresholding of the iterates. In the matrix case, the proposed strategy yields iterates whose ranks remain comparable to natural benchmark quantities — namely, the best approximation ranks of the sought solutions at the achieved accuracy. In the higher-dimensional tensor case the algorithm can be adapted appropriately, leading to global

error and rank bounds that depend only polynomially on the dimension. Numerical experiments illustrate the theory for linear time-dependent Schrödinger equations.

Adaptive Randomized Tensor Train Rounding using Khatri-Rao Products

Authors: Arvind Krishna Saibaba[1], Bhisham Dev Verma[2], Grey Ballard[2], Hussam Al Daas[3], Laura Grigori[4], Mariana Martinez Aguilar[5]

[1] *North Carolina State University*

[2] *Wake Forest University*

[3] *Scientific Computing Department, STFC, Rutherford Appleton Laboratory, UK*

[4] *Institute of Mathematics, EPFL, Lausanne, and PSI Center for Scientific Computing, Theory and Data, Villigen PSI, Switzerland*

[5] *Institute of Mathematics, EPFL, Lausanne, Switzerland*

Contact: vermabd@wfu.edu

The Tensor Train (TT) format provides a compact and scalable way to represent high-dimensional tensors, making it essential for solving certain parametrized partial differential equations and other large-scale problems. A critical operation in TT-based computations is rounding, which reduces the ranks of a tensor in TT format to maintain efficiency. While recent randomized rounding algorithms have improved performance over traditional deterministic approaches, they typically require manual rank specification, limiting their practicality in many applications.

In this talk, I will present a new adaptive randomized TT-rounding algorithm based on Khatri-Rao random projections. Unlike previous methods, our algorithm automatically determines the TT ranks needed to meet a user-specified approximation tolerance, eliminating the need for manual tuning. It achieves up to 45× speedup over deterministic rounding and a 2× improvement over the fastest existing randomized approach (which requires specifying target ranks), particularly in the context of rounding sums of TT tensors—a key bottleneck in adaptation of GMRES to vectors in TT format. I will discuss the algorithm's design, performance benefits, and its implications for high-dimensional numerical computing.

Decomposition of a tensor into multilinear rank-$(M_r, N_r, \cdot)$ terms

Authors: Nico Vervliet[1], Lieven De Lathauwer[1]

[1] *KU Leuven*

Contact: nico.vervliet@esat.kuleuven.be

We present new generic and deterministic uniqueness results for block term decompositions (BTD). These uniqueness conditions hold under mild assumptions and apply to more general settings than previously known results. We also present an algebraic algorithm for the computation of BTDs. Our algorithm requires no knowledge of the block sizes appearing in the BTD: these block sizes are recovered from the algorithm. Through numerical simulations, we illustrate that, in contrast to competing optimization-based methods, even in noisy settings our algebraic algorithm can successfully recover an underlying BTD without knowledge of block sizes provided the signal-to-noise ratio is sufficiently high. We observe that the algorithm can significantly improve one's ability to successfully recover a BTD when it is used as an algebraic initialization for leading optimization routines. Moreover, only a few optimization iterations are required to successfully converge to the BTD from the algebraic solution.

Low-Rank CP Tensor Compression and Its Application to High-Dimensional PDEs

Authors: Zhanrui Zhang[1], Edgar Solomonik[1]

[1] *University of Illinois Urbana-Champaign*

Contact: zhanrui3@illinois.edu

The Canonical Polyadic (CP) decomposition is widely used to represent high-dimensional data in many applications, for example, solving high-dimensional PDEs like kinetic equations. A key challenge in these problems is the efficient estimation and reduction of the CP rank. The CP rank reduction task can be formulated as approximating the Khatri–Rao product of the CP factor matrices with a lower rank. We propose an approach based on the pivoted Cholesky decomposition to construct interpolative decompositions of the Khatri–Rao product. This method can serve as a standalone CP rank reduction technique or be integrated into classical optimization schemes such as CP-ALS. Moreover, the residuals produced at each step of the pivoted Cholesky naturally provide error indicators, enabling effective rank estimation. Preliminary numerical results demonstrate that this method achieves a balance between computational cost and approximation accuracy. Its effectiveness is further validated through applications to the Vlasov–Poisson equation.

MS 32: Advances in Randomized Algorithms and Kernel Methods for Rank-structured Matrices

Minisymposium Organizers: Mikhail Lepilov, Andrew Horning

Rank-structured matrices arise in many areas of applied mathematics, including data science, machine learning, particle systems, and numerical methods for partial differential equations. This minisymposium aims to bring together researchers working in two active areas of rank-structured linear algebra: randomized numerical linear algebra and low-rank techniques for kernel matrices. The speakers will present new results and algorithms for spectral properties of rank-structured matrices, as well as new and improved techniques—leveraging analytic, algebraic, and randomized methods—for computing their low-rank decompositions. Examples of topics include randomized methods for density-of-states, as well as the proxy point method for analytic kernel approximation. Domain-specific applications in data-science and large-scale scientific simulations will also be presented.

List of Talks

Eigenvector Approximation via Random Sampling

Authors: Cameron Musco, Rajarshi Bhattacharjee[1]

[1] *University of Massachusetts Amherst*

Contact: rbhattacharj@umass.edu

We study the problem of approximating the eigenvectors of an n x n symmetric matrix A with bounded entries using random column sampling. We show that for any eigenvalue λ of A, one can compute a vector v satisfying $||Av - \lambda v||_2 \leq \epsilon n$ by sampling $\tilde{O}(\frac{1}{\epsilon^4})$ columns of A. For the eigenvector corresponding to the largest-magnitude eigenvalue, this sample complexity improves to $\tilde{O}(\frac{1}{\epsilon^2})$ columns.

We further establish a lower bound on the number of entries of A that must be sampled to approximate its top eigenvector up to error ϵn, showing that our algorithm is optimal up to polylogarithmic factors. When columns of A can be sampled with probabilities proportional to their squared l_2-norms, we obtain an improved error guarantee of $\epsilon||A||_F$ for the approximate eigenvectors, where $||A||_F$ denotes the Frobenius norm.

Our analysis builds on recent results of Bhattacharjee et al. (2022) and Swartworth and Woodruff (2025), which demonstrate that all eigenvalues of A can be approximated up to additive error ϵn (and $\epsilon||A||_F$ under l_2 sampling) by sampling a sufficiently large principal submatrix.

Acceleration of Lanczos approximation for PDE discretizations in unbounded domains

Authors: Vladimir Druskin[1], Jorn Zimmerling[2]

[1] *Southern Methodist University*

[2] *Uppsala University*

Contact: vdruskin@gmail.com

We consider the approximation of $B^T(A + sI)^{-1}B$ for large s.p.d. $A \in \mathbb{R}^{n\times n}$ with a dense spectrum and $B \in \mathbb{R}^{n\times p}$, $p \ll n$ using block-Lanczos recursion. We target the computations of MIMO transfer functions for large-scale discretizations of problems with continuous spectral measures, such as linear time-invariant (LTI) PDEs on unbounded domains. While Krylov methods such as Lanczos and CG are near-optimal for problems with discrete, well-separated spectra, their spectral adaptation deteriorates in the dense-spectrum regime. We address this limitation using a square-root terminator framework (originated in quantum physics in 1970s), modifying the final Lanczos recursion coefficient in an s-dependent manner via a Krein–Nudelman damped string representation. The resulting approximants reproduce the same Stieltjes moments while yielding continuous spectral measures and significantly reduced approximation error through adaptive damping, that maximizes the relative outflow of energy.

Large-scale experiments for diffusion and wave PDEs demonstrate the competitiveness of the proposed approach for deterministic as well as randomized computations.

Low-Rank Approximation by Randomly Pivoted LU

Authors: Heather Wilber[1], Marc Aurèle Gilles[2]

[1] *University of Washington*

[2] *Princeton University*

Contact: gilles@princeton.edu

I will present Randomly Pivoted LU (RPLU), a randomized variant of Gaussian elimination with complete pivoting that samples pivots proportional to squared Schur-complement entries, and analyze its low-rank approximation properties. I will highlight two regimes where RPLU is particularly effective at low-rank approximation: (i) memory-limited settings, where a rank-k approximation can be computed with $\mathcal{O}(k^2 + m + n)$ storage and $\mathcal{O}(k^3 + m + n + k\mathcal{M}(\mathbf{A}))$ work (with $\mathcal{M}(\mathbf{A})$ the cost of a matvec with $\mathbf{A} \in \mathbb{C}^{m\times n}$ or $\mathbf{A}^*$), and (ii) structured problems where $\mathbf{A}$ and its Schur complements admit fast updates (e.g., Cauchy-like matrices). I will discuss applications to fast computation of high-degree rational approximants.

Proxy Points and Rational Approximation Using Contour Integration

Authors: Mikhail Lepilov[1], Andrew Horning[1]

[1] *Rensselaer Polytechnic Institute*

Contact: mlepilov@gmail.com

Due to the size of many kernel matrices that arise in applications, it is often necessary to work with their low-rank approximations in order to efficiently perform many computations. Low-rank matrix decompositions to such matrices may be quickly obtained by exploiting the analytic structure of the underlying kernel, for example by using Taylor expansions or an integral representation; such ideas trace back to the fast multipole method. However, there is often a gap between the theoretical epsilon-rank of a given kernel matrix, explicitly obtained using its eigenvalue decomposition, and the existing analytic decompositions of the same rank. In this work, we aim to bridge this gap by exploring the proxy point method, which approximates complex-analytic kernel matrices using a discretized contour integral representation. In particular, we combine this approach with a judicious choice of conformal map applied to the set of points at which the kernel is evaluated in order to obtain effective analytic approximations. We apply this new method to some well-studied kernel matrices, such as the Hilbert and Cauchy matrices, in order to compare it with existing analytic methods and with the best available theoretical bounds. In the process, we demonstrate a novel way of obtaining one kind of theoretically-optimal solution to a related problem in rational approximation.

Kernel-Based Variational Formulations of Nonlocal Mean-Field Games

Author: Siting Liu[1]

[1] *University of California, Riverside*

Contact: sitingl@ucr.edu

We study mean-field games with nonlocal interactions modeled through kernel-based representations. Using feature-space expansions inspired by kernel methods, the resulting models admit a variational saddle-point formulation that is well suited for efficient primal–dual algorithms such as the primal-dual hybrid gradient method. We also discuss inverse problems in which interaction mechanisms are recovered from partial and noisy observations of population dynamics, emphasizing optimization-based formulations and splitting methods for stable and efficient reconstruction.

Practical Spectral Density Estimation with Explicit Deflation

Authors: Ibrohim Nosirov[1], Anil Damle[1], Apoorv Vikram Singh[2], Christopher Musco[2]

[1] *Cornell University*

[2] *New York University*

Contact: in67@cornell.edu

Stochastic Lanczos Quadrature (SLQ) is a popular algorithm for approximating the spectral density of a symmetric matrix A using matrix-vector products. We present a variance reduced implementation of SLQ. This implementation has two key ingredients: a faster problem-specific eigensolver and a carefully implemented selective orthogonalization scheme that we use as a deflation criterion. Our eigensolver is observed to be faster, more robust, and to scale better than LAPACK's 'stemr' (MRRR) in the context of SLQ. Equipped with this faster eigensolver, we explicitly track residual information and perform deflation to speed up convergence. This is achieved using an implementation that closely follows the LanSO algorithm described in Parlett's *The Symmetric Eigenvalue Problem*.

MS 33: Quantum Numerical Linear Algebra

Minisymposium Organizers: Daan Camps, Arielle Carr, Kathryn Lund, Roel Van Beeumen

Quantum computing is rapidly emerging as a disruptive technology throughout the sciences and industry. Hardware improvements have led to an increased interest in discovering new quantum algorithms that can deliver on the promise of quantum advantage. Recent progress has made it increasingly clear that quantum linear algebra kernels are at the core of many promising quantum algorithms in material science, chemistry, condensed matter physics, and machine learning. At the same time, quantum computers introduce new sources of error at both the hardware and algorithmic levels. It is therefore imperative that classically trained numerical linear algebraists become involved in the development of new quantum algorithms and technology. This minisymposium covers a broad range of novel advances in quantum linear algebra methods and their applications in the sciences. Examples of algorithms include but are not limited to quantum signal processing, quantum singular value transformation, quantum subspace methods, variational algorithms, quantum linear systems, and quantum eigenvalue problems. Application areas of interest include but are not limited to quantum machine learning, computational chemistry, and condensed matter physics. We aim to provoke discussion on key open questions in the field and highlight the potential for classical numerical linear algebraic solutions to quantum problems.

List of Talks

Efficient Encoding of Semiseparable Matrices in Quantum Circuits

Authors: Gianna M. Del Corso[1], Giacomo Antonioli, Roel Van Beeumen

[1] *Dip. Informtica, Università di Pisa*

Contact: gianna.delcorso@unipi.it

Quantum block encoding (QBE) is a crucial step in the development of most quantum algorithms, providing an embedding of a given matrix into a suitable larger unitary matrix. Efficient techniques for QBE have primarily focused on sparse matrices, and less effort has been devoted to data-sparse matrices, such as rank-structured matrices.

In this talk, we examine a specific case of rank structure: one-pair semiseparable matrices. We present a new block encoding approach that relies on an efficient state preparation technique for the representation of the two generators and on the reordering of entries obtained through the tensor product of the generators. This process takes polylogarithmic time if the two generators can be prepared efficiently, with a scaling factor of $2\sqrt{2}^n$, where $n = log_2(N)$ is the number of qubits to represent the matrix.

Practical block encodings of matrix polynomials that can also be trivially controlled

Authors: Filippo Della Chiara[1], Martina Nibbi[2], Roel Van Beeumen, Yizhi Shen[3], Aaron Szasz[4]

[1] *KU Leuven*

[2] *Technical University of Munich*

[3] *Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory*

[4] *Google Quantum AI*

Contact: filippo.dellachiara@kuleuven.be

Quantum circuits naturally implement unitary operations on input quantum states. However, non-unitary operations can also be implemented through "block encodings", where additional ancilla qubits are introduced and later measured. While block encoding has a number of well-established theoretical applications, its practical implementation has been prohibitively expensive for current quantum hardware. In this paper, we present practical and explicit block encoding circuits implementing matrix polynomial transformations of a target matrix. With standard approaches, blockencoding a degree-d matrix polynomial requires a circuit depth scaling as d times the depth for block-encoding the original matrix alone. By leveraging the recently introduced Fast One-Qubit Controlled Select LCU (FOQCS-LCU) framework, we show that the additional circuit-depth overhead required for encoding matrix polynomials can be reduced to scale linearly in d with no dependence on system size or the cost of block encoding the original matrix. Moreover, we demonstrate that the FOQCS-LCU circuits and their associated matrix polynomial transformations can be controlled with negligible overhead, enabling efficient applications such as Hadamard tests. Finally, we provide explicit circuits for representative spin models, together with detailed non-asymptotic gate counts and circuit depths.

Computational linear algebra from time evolution and noisy inner products

Author: Ethan Epperly[1]

[1] *UC Berkeley*

Contact: eepperly@berkeley.edu

The fundamental building blocks of iterative linear algebra algorithms in ordinary digital computation are matrix–vector multiplications and inner products. In quantum computing, we lose easy access to both primitives. But we also gain replacements. Instead of matrix–vector products, we can apply the unitary time evolution operator e^{-itA}, and we have access to noisy—but statistically unbiased—measurements of inner products. How can we solve linear algebra problems with these primitives? What is the greatest possible accuracy achievable? This talk will address these questions with a focus on eigenvalue problems and on the presenter's own research. No prior knowledge of quantum computing is expected.

Quantum Krylov methods with Hamiltonian powers

Author: Ryan LaRose
Contact: rmlarose@msu.edu

Quantum Krylov methods are strong candidates for computing ground states on NISQ and Mega-Quop computers. While typically implemented with powers of the time evolution unitary e^{-iHt} for a given Hamiltonian H, convergence can be markedly faster with powers of the Hamiltonian H itself as in classical methods. We discuss these convergence rates and present several ways to implement Hamiltonian powers on quantum computers, including qubitization and randomized methods, for use in quantum Krylov. We compare the approaches by providing resource estimates for common applications on current and future quantum computers.

Quantum Linear Algebra for Optmization

Author: Mohammadhossein Mohammadisiahroudi[1]

[1] *Department of Mathematics and Statistics, University of Maryland Baltimore County*

Contact: mhms379@umbc.edu

Quantum linear algebra has emerged as a promising framework for accelerating the solution of fundamental computational problems, including systems of linear equations—a core subroutine in many scientific and engineering tasks. These problems arise prominently in optimization algorithms. In this talk, we discuss the opportunities and challenges associated with integrating quantum linear algebra techniques into modern optimization methods. We focus on two key applications:

quantum interior point methods for conic optimization and a quantum adjoint method for PDE-constrained optimization. Together, these illustrate how quantum algorithms may reshape large-scale optimization by offering new pathways toward improved scalability and performance.

Quantum Matrix Encodings

Authors: Liron Mor Yosef[1], Haim Avron[1]

[1] *Tel Aviv University*

Contact: lm2@mail.tau.ac.il

The emergence of Quantum Numerical Linear Algebra (qNLA) offers a paradigm shift in solving large-scale linear systems and matrix functions. However, the practical utility of these algorithms, such as the seminal HHL, is fundamentally bottlenecked by the "input problem", namely the efficient representation of classical matrices as quantum circuits.

In this talk, we explore two distinct methodologies for addressing this challenge. First, we present a systematic framework for constructing block encodings from state preparation circuits. This is achieved by introducing efficient basis-transformation protocols and low-overhead, constant-depth Pauli multiplexers. We then establish a bidirectional operational equivalence by demonstrating the converse: converting block encodings back into state preparation circuits using quantum linear algebra operations.

Furthermore, we introduce a structurally-driven approach based on the factorization of matrix state preparation. We prove that when the state preparation algorithm is applied to vectorized classical data, this algorithm naturally decomposes into a standard Norm Loader and a Direction Multiplexer. This observation reveals that a well-known approach for creating block encodings was effectively "hidden" within the classical logic of state preparation. By leveraging this hidden structure, we present a new method for block encoding that offers unique advantages in qubit utilization and circuit depth.

Efficient LCU block encodings through Dicke states preparation

Authors: Filippo Della Chiara[1], Martina Nibbi[2], Yizhi Shen[3], Roel Van Beeumen[3]

[1] *KU Leuven*

[2] *Technical University of Munich*

[3] *Lawrence Berkeley National Laboratory*

Contact: rvanbeeumen@lbl.gov

With the Quantum Singular Value Transformation (QSVT) emerging as a unifying framework for diverse quantum speedups, the efficient construction of block encodings—their fundamental input

model—has become increasingly crucial. However, devising explicit block encoding circuits remains a significant challenge. A widely adopted strategy is the Linear Combination of Unitaries (LCU) method. While general, its practical utility is often limited by substantial gate overhead. To address this, we introduce the Fast One-Qubit-Controlled Select LCU (FOQCS-LCU), a compact LCU formulation that requires only a linear number of ancilla qubits and is explicitly decomposed into one- and two-qubit gates. By exploiting the underlying Hamiltonian structure, we design a parametrized family of efficient Dicke state preparation routines, enabling systematic realization of the state preparation oracle at substantially reduced gate cost. The check matrix formalism further yields a constant-depth SELECT oracle, implemented as two fully parallelizable layers of singly controlled Pauli gates. We construct explicit block encoding circuits for representative spin models such as the Heisenberg and spin glass Hamiltonians and provide detailed, non-asymptotic gate counts. Our numerical benchmarks confirm the efficiency of the FOQCS-LCU approach, illustrating over an order-of-magnitude reduction in CNOT count compared to conventional LCU. This framework opens a pathway toward practical, low-depth block encodings for a broad class of structured matrices beyond those considered here.

Quantum Linear Algebra: from Optimization to Differential Equation

Authors: Zeguan Wu[1], Marc Vuffrey[2], Xiu Yang[3], Tamás Terlaky[3], Junyu Liu[1]

[1] *University of Pittsburgh*

[2] *Los Alamos National Laboratory*

[3] *Lehigh University*

Contact: zew79@pitt.edu

Quantum computing relies heavily on the efficient manipulation of linear algebraic structures. This talk discusses the application of quantum linear algebra across two major domains: continuous optimization and differential equations. We demonstrate how quantum linear algebra can be utilized to solve these problems efficiently, discussing both the algorithmic construction and the theoretical analysis required to establish quantum advantage.

MS 34: Model- and Data-driven Reduced-order Models and Their Applications in Inverse Problems

Minisymposium Organizers: Vladimir Druskin, Mikhail Zaslavskiy, Joern Zimmerling

This minisymposium aims to explore recent advancements in model- and data-driven reduced order models (ROMs) and their applications in imaging and inverse scattering. The ROMs have been shown to be a powerful tool for accelerating the solution of large-scale multi-dimensional inverse problems, however many challenges remain. They include: Construction of structure-preserving ROMs; Design of efficient linear algebraic algorithms for large datasets; Computation of transfer functions for problems with almost-continuous spectra, Truncation and regularization of data-driven Gramians.

List of Talks

Characterization of the response of electric circuits with two kinds of passive elements

Authors: Fernando Guevara Vasquez[1], Graeme Milton, Aaron Welters

[1] *University of Utah*

Contact: fguevara@math.utah.edu

The response matrix of an electrical circuits composed of two kinds of passive elements (e.g. LC, RC, CL...) is a matrix-valued rational function of the frequency that associates time-harmonic voltages at some terminal nodes to time-harmonic currents. We present necessary conditions that must be satisfied by the response matrix. These conditions can be used to construct circuits whose

response matches prescribed poles, with some exceptions. Progress towards a full characterization will be presented.

Multiparameter Waveform Inversion via Reduced Order Modeling

Authors: Alexander Mamonov[1], Josselin Garnier, Jörn Zimmerling[2], Liliana Borcea

[1] *University of Houston*

[2] *Uppsala University*

Contact: mamonov@gmail.com

Waveform inversion seeks to estimate an inaccessible heterogeneous medium from data gathered by sensors that emit probing signals and measure the generated waves. The traditional full waveform inversion (FWI) formulation estimates the unknown coefficients via minimization of the nonlinear, least squares data fitting objective function. For typical band-limited and high frequency data, this objective function has spurious local minima near and far from the true coefficients. Thus, FWI implemented with gradient based optimization algorithms may fail, even for good initial guesses, a phenomenon known as cycle skipping. Recently, it was shown that data driven reduced order models (ROMs) can be used to obtain a better behaved objective function for wave speed estimation. We introduce ROMs for waves obeying a first order hyperbolic system. They are defined via Galerkin projection on the space spanned by the wave snapshots, evaluated on a uniform time grid with an appropriately chosen time step. The proposed ROMs are data driven, as they are computed directly from the sensor measurements. The ROM computation applies to any linear waves in lossless and non-dispersive media. We present an example of acoustic waveform inversion in a medium with both wave speed and density unknown. Numerical examples show that both quantities can be estimated efficiently via minimization of an objective function that employs ROM based approximation of the wave field inside the unknown medium.

Reduced order models for inverse scattering

Authors: Alexander Mamonov, Andreas Tataris
Contact: andreas.tataris@ensta.fr

We present a numerical method for solving an inverse boundary value problem of estimating the acoustic velocity in the Helmholtz equation from frequency domain measurements based on reduced order models (ROM). The ROM is the Galerkin projection of the Helmholtz operator onto a subspace spanned by its solution snapshots at certain wavenumbers. We show how to reconstruct the ROM in a data-driven way. Once the ROM is reconstructed, the acoustic velocity can be estimated using non-linear optimization that minimizes a misfit based on the ROM. Such an approach

typically outperforms the conventional methods based on data misfit minimization.

Adaptive data-driven reduced-order models of port-Hamiltonian dynamical systems for nonlinear inverse scattering applications

Author: Mikhail Zaslavskiy[1]

[1] *Southern Methodist University*

Contact: mzaslavskiy@smu.edu

The inverse scattering problem formulated for the Schrödinger operators arises in various fields, including quantum mechanics, radars, viscoelasticity, Biot problems, remote sensing, geophysical, and medical imaging. The goal of imaging is to find medium properties in the domain using near-field measured data. The model based nonlinear optimization which is the method of choice for the solution of the inverse problems can be unreliable and particularly expensive for such problems. Data driven nonlinear transforms can be an opening, however it was recently shown that the ReLU networks are intractable for reliable solution of the inverse problems in continuum using conventional digital computers. In the present work, following the success of data-driven reduced-order models (ROMs) developed in recent years, we propose a robust direct method for solving inverse scattering problems for the Schrödinger equation. Our approach is based on a Lippmann-Schwinger algorithm with a crucial component composed of adaptive data-driven ROMs in the frequency domain and efficient learning the internal solutions. In the talk we will discuss the details of the algorithms as well as some bottlenecks.

ROM-based Inverse Scattering for Monostatic Data

Author: Jörn Zimmerling[1]

[1] *Uppsala University*

Contact: jorn.zimmerling@it.uu.se

We consider an inverse scattering problem for monostatic synthetic aperture radar (SAR), where the goal is to estimate the wave speed in a heterogeneous, isotropic medium using measurements from a moving antenna. The forward map, derived from Maxwell's equations, is inherently nonlinear, accounts for multiple scattering, and exhibits high-frequency oscillations that make traditional nonlinear least-squares formulations prone to local minima.

We present an alternative two-step approach centered on the construction of data-driven Reduced Order Models (ROMs). The first step computes a non-iterative map from the measurements to an approximation of the internal electric field. This internal wave is designed to fit the data by construction. The second step utilizes an iterative optimization to minimize the discrepancy

between this internal field and the solution to the governing Maxwell's equations across all antenna locations.

While the ROM-based approach provides a superior estimate of the target support compared to standard SAR imaging, its performance is sensitive to the algebraic properties of the data-driven operators. In this talk, we focus on the linear algebraic challenges involved in stabilizing these ROMs. Specifically, we discuss the spectral properties of the data-driven mass and stiffness matrices and review regularization techniques to maintain robustness when the monostatic data is polluted with noise and preserve the structure of the ROM. Numerical simulations demonstrate that these stabilized ROMs allow for accurate quantitative estimation of the wave speed in regimes where standard inversion methods fail.

MS 35: Recent Advances in Tensor Decompositions for Model and Data Reduction

Minisymposium Organizers: Misha Kilmer, Vishwas Rao

Modern scientific applications generate increasingly large datasets. These datasets are typically stored as multi-dimensional arrays whose entries correspond to values of physical quantities in spatio-temporal coordinates. Traditional matrix-based reduction techniques struggle to capture the multiway correlations inherent in such data, resulting either in excessive loss of structure or prohibitive computational cost. Tensor-based data reduction offers a principled and scalable alternative by representing high-dimensional datasets through structured multi-linear factorizations such as CANDECOMP–PARAFAC (CP), starM, Tucker, Tensor Train (TT), and Hierarchical Tucker (HT) decompositions. These approaches exploit low-rank structure across multiple modes simultaneously, enabling compression, noise filtering, and feature extraction far beyond what is possible with classical methods. This minisymposium will highlight recent methodological and computational advances in tensor decompositions for streaming data and models, surrogate models, and high-performance implementations of these decompositions.

List of Talks

Derivative informed Tucker tensor train Taylor series surrogate models

Authors: Blake Christierson[1], Nick Alger[1], Omar Ghattas[1], Peng Chen[2]

[1] *University of Texas, Austin*

[2] *Georgia Tech*

Contact: nalger225@gmail.com

We introduce Tucker tensor train Taylor series (T4S) surrogate models for high dimensional mappings that depend implicitly on the solution of a partial differential equation. Traditionally, Taylor series are intractable here because the derivative tensors are enormous, and are only accessible through multilinear actions. We overcome these challenges by approximating each derivative tensor with a Tucker decomposition composed with a tensor train, fitting each Tucker tensor train to symmetric tensor actions via Riemannian manifold optimization. We present theory and numerical experiments that validate the model and numerical method.

Streaming Compression of Scientific Data through Weak-SINDy and POD Integration

Author: Rick Archibald[1]

[1] *Oak Ridge National Laboratory*

Contact: archibaldrk@ornl.gov

The exponential growth of scientific data from simulations and experiments demands efficient compression techniques for storage and processing. This talk introduces a novel streaming weak-SINDy algorithm designed for real-time compression of streaming scientific data. Leveraging the underlying structure of physical systems, the algorithm constructs memory-efficient feature matrices and target vectors in real-time, enabling model recovery in an offline regression stage. For high-dimensional data, we integrate a streaming proper orthogonal decomposition (POD) process, dynamically reducing data dimensions and augmenting the streaming weak-SINDy algorithm to handle evolving POD bases. Proof-of-concept examples, including applications to the Lorenz system and fluid-flow data, demonstrate both the memory efficiency and reconstruction accuracy of the approach. Join us to explore how these advancements pave the way for scalable, real-time compression of scientific data.

High-Performance Implementation of Star-M SVD for Big Data Compression

Author: Md Taufique Hussain[1]

[1] *Wake Forest University*

Contact: hussaint@wfu.edu

In the era of big data, effectively compressing large datasets while performing complex mathematical operations is crucial. Tensor-based decomposition methods have shown superior compression capabilities with minimal loss of accuracy compared to traditional matrix methods. Under the

$\star_M$ tensor framework, tensors can be decomposed in a matrix-mimetic way, including using the $\star_M$ SVD. This tensor SVD has optimality guarantees and has shown exceptional performance on specific types of data, but software implementations have been mostly limited to productivity-oriented languages. In this work, we present our development of a shared-memory parallel, high-performance solution designed to efficiently implement the underlying algorithms. This software will enable optimal compression of extensive scientific datasets, paving the way for enhanced data analysis and insights.

Scalable Moment Tensor Decompositions

Author: Joe Kileel[1]

[1] *University of Texas, Austin*

Contact: jkileel@math.utexas.edu

Low-rank decompositions of moment tensors provide a natural approach for density and parameter estimation of mixture models, given data samples. In general, the mixture condition corresponds to the structure of a convex linear combination in terms of moment tensors. Therefore, moments of mixture models have low-rank tensor decompositions. Furthermore, computing the low-rank factors often reveals the parameters of interest quite directly.

All of this is rather well known in theory. But thus far, it has had a muted impact on practice, primarily because computations with high-order and high-dimensional tensors seem prohibitively expensive.

In this talk, I will describe work I have been doing that breaks down this barrier. The main theme will be the development of numerical algorithms for moment tensor decompositions that evade the apparent curse of dimensionality inherent in moment tensors. I will discuss results for a range of mixture models, including Gaussian mixture models, mixtures of products, and mixtures of nonparametric distributions with banded correlations. Time permitting, some real-life scientific applications will be mentioned, to highlight the practical improvements brought by the new methods.

Improving prediction for a low rank tensor ROM via continuous data assimilation

Author: Leo Rebholz[1]

[1] *Clemson University*

Contact: rebholz@clemson.edu

We extend a low rank tensor ROM recently developed by Olshanskii et al by enhancing it with continuous data assimilation (CDA). We show how CDA is easily incorporated into the ROM, and analytically show that it provides for theoretical long time error estimates. Numerical tests illustrate the theory and show it is an effective tool for simulating incompressible flow over a wide range of Reynolds numbers.

POD-DEIM in the starM-product framework

Authors: Amit Subrahmanya[1], Arvind Krishna Saibaba[2], Eric de Sturler[3], Misha Kilmer[4], Vishwas Rao

[1] *Argonne National Laboratory*

[2] *North Carolina State University*

[3] *Virginia Tech*

[4] *Tufts University*

Contact: vhebbur@anl.gov

We propose using the starM tensor product framework for constructing Proper Orthogonal Decomposition (POD) and Discrete Empirical Interpolation Method (DEIM) reduced order models. By exploiting the inherent multidimensional relationship structure of snapshot data, the approach enables efficient computation of the reduced bases. Operating directly on tensor representations reduces storage and computational costs while maintaining accuracy of the full-order dynamics compared to the more standard vector/matrix based methods. The resulting *M POD–DEIM models are well suited for large-scale, high-dimensional systems, and we demonstrate it on a number of different test problems.

Streaming Tensor BM-Decomposition

Author: Fan Tian
Contact: fantian93@gmail.com

Tensor decomposition is widely used for analyzing multi-way data in various applications that often involve continuously generated data. Efficient methods to process and extract meaningful patterns dynamically are hence essential for these applications. In this talk, we consider the problem of computing the streaming tensor BM-decompositions (BMD). An incremental algorithm, OnlineBMD, is presented, which is the first work that addresses third-order tensor BMD in a dynamic setting.

MS 36: Polynomials, Krylov Methods and Applications

Minisymposium Organizers: Cade Ballew, Ethan Epperly, Alexander Hsu

For many large-scale problems in numerical linear algebra, iterative methods are the most effective and sometimes the only tractable approach. The most common approaches approximate a desired quantity in the Krylov space, which is intimately related to polynomial approximation. In the era of big data and increasing computational demands, there has been a flurry of interest in such iterative methods and efficient preconditioners. This minisymposium will bring together researchers across fields who analyze, develop, and apply Krylov and polynomial methods interpreted broadly.

List of Talks

Spectral density estimation for random matrices

Authors: Charbel Abi Younes[1], Thomas Trogdon[1], Xiucai Ding[2]

[1] *University of Washington*

[2] *UC Davis*

Contact: cyounes@uw.edu

We introduce a new approach for estimating the asymptotic spectral distribution (ASD) of a random matrix using a single, sufficiently high-dimensional sample, without computing the full spectrum. The method builds on the Lanczos algorithm, together with asymptotic analysis and perturbation theory for orthogonal polynomials, and enables efficient and accurate estimation of the ASD. We illustrate the approach through an application to spectral density estimation in spiked covariance models.

Computing Functions of Rank-structured or Telescopic Matrices

Authors: Andrea Baleani[1], Leonardo Robol[2], Stefano Massei[2]

[1] *Scuola Normale Superiore di Pisa*

[2] *University of Pisa*

Contact: andrea.baleani@sns.it

Appearing in a wide variety of applications, often in the context of discretized (fractional) differential and integral operators, Hierarchically Semiseparable (HSS) matrices have a number of attractive properties facilitating the development of fast algorithms [6,4].

For HSS matrices, the rank-structure is numerically preserved if $f(z)$ is well-approximated by a rational function. The underlying reason for this is that if A is HSS, A^{-1} shares the same structure [5]. Recently, Casulli, Kressner and Robol proposed an algorithm for computing matrix functions $f(A)$ where A is HSS and symmetric [3]. This result is based on the equivalence between HSS structure and the existence of telescopic decompositions

$$A = A^{(L)} = D^{(L)} + U^{(L)} A^{(L-1)} V^{(L)^T}.$$

In our work we extend the procedure to general (non-symmetric) HSS matrices. The algorithm works recursively, viewing the telescopic decomposition as a sequence of low-rank updates to block-diagonal matrices $D^{(L-i)}$ for $i = 0, \dots, L$. At each step, after some careful permutation of the diagonal blocks, we employ Block Rational Krylov methods [1] to approximate the effect of the low-rank update while keeping the size of the matrices involved small.

Moreover, this technique allows the efficient computation of $f(T)$ where T is a Toeplitz. Up to an FFT, the off-diagonal blocks of a Toeplitz matrix satisfy a displacement equation. The low-rank of the right-hand side is linked to fast singular value decay, which, in turn, implies that $\Omega T \Omega^*$ can be well-approximated by an HSS matrix [2].

[1] B. Beckermann et al., SIAM J. Numer. Anal., 59(3), 2021.
[2] B. Beckermann et al., Numerical Algorithms 100(4), 2025.
[3] A. A. Casulli et al., SIAM J. Matrix Anal. Appl., 45(4), 2024.
[4] S. Chandrasekaran et al., SIAM J. Matrix Anal. Appl., 27(02), 2005.

[5] A. Cortinovis et al., SIAM J. Matrix Anal. Appl., 43(1), 2022.
[6] W. Hackbusch, *Hierarchical Matrices: Algorithms and Analysis*, 49, 2015.

Tensor Network Krylov Methods: Algorithms, Theory Gaps, and Open Problems

Author: Chris Camaño[1]

[1] *Caltech*

Contact: ccamano@caltech.edu

In recent years, *tensor network methods* have garnered increased attention for modeling high dimensional quantum many body systems and for representing high dimensional functions with structured correlations.

A basic unresolved question is whether standard numerical linear algebra primitives, such as Krylov based iterative solvers, can be realized in matrix product state and matrix product operator form while remaining computationally efficient.

This talk gives a pedagogical overview of the current state of tensor network Krylov methods. I will define what it means to run a Krylov iteration when vectors and operators are represented in the tensor network format, and explain why familiar steps like basis construction and reorthogonalization become nontrivial algorithmic design problems in this setting. I will then show how tensor network structure can be paired with randomized low rank approximation to speed up the tensor network analog of matrix vector products, opening a broader design space for fast high dimensional Krylov solvers. Recent approaches will be surveyed, including tensor network Lanczos variants and sketched tensor network GMRES methods.

I will close with open questions around stability, error control, and tensor network bond dimension growth that are currently shaping the development of new algorithms in this space.

Randomized Row Norm Estimation: Algorithms and Applications

Authors: Alexander Hsu[1], Ethan Epperly[2]

[1] *University of Washington*

[2] *UC Berkeley*

Contact: owlx@uw.edu

Computing the diagonal entries of a large linear operator is a common computational primitive in numerical linear algebra, with applications in uncertainty quantification, cross-validation, perturbation analysis, electronic structure calculation and more. However, estimating the diagonals of a

matrix given only implicit matrix-vector access is challenging, as randomized algorithms suffer from high variance. In many cases, these problems can be efficiently reduced to estimating the squared row norms of a suitable square root of the matrix, where randomized algorithms specialized to row norm estimation with considerably lower variance may be applied. In particular, when the matrix in question is a function of another matrix, that is, $A = f(B)$, and matrix vector products with A are computed using a Krylov method, it is often the case that matrix vector products with $f^{1/2}(B)$ can be computed just as efficiently those with $f(B)$. We present algorithms for estimating squared row norms and approaches for incorporating them into other computational tasks.

Stochastic trace estimation for parameter-dependent matrices

Authors: Fabio Matti[1], Haoze He, Daniel Kressner[1], Hei Yin Lam

[1] *EPFL*

Contact: fabio.matti@epfl.ch

Stochastic trace estimators are a family of widely used techniques for approximating traces of large matrices accessible only via matrix-vector products. These methods have been studied extensively when applied to *constant* matrices B. We analyze three standard stochastic trace estimators—the Girard-Hutchinson, Nyström, and Nyström++ estimators—when they are applied to *parameter-dependent* matrices $B(t)$ that continuously depend on a real parameter $t \in [a, b]$. Our key observation is that a single set of random vectors can be reused to form the estimators for all values of t, yielding estimates whose L^1-error bounds match those of the constant-matrix case. Traces of parameter-dependent matrices arise naturally in important applications, including spectral density estimation and partition function estimation. Building on our analysis, we develop algorithms that combine Chebyshev interpolation with parameter-dependent stochastic trace estimation to obtain efficient methods with provable accuracy guarantees for these problems.

The Matrix-Vector Complexity of Ax=b

Authors: Ethan Epperly[1], Michał Dereziński[2], Raphael Meyer[1]

[1] *UC Berkeley*

[2] *University of Michigan*

Contact: ram900@berkeley.edu

Solving linear systems Ax=b is a fundamental pillar of NLA. For over 60 years, iterative methods that access A only through matrix-vector products have been the standard approach for solving large linear systems. While lower bounds exist for many special cases, prior work has not shown that methods like GMRES and MINRES achieve an asymptotically optimal matrix-vector complexity for approximately solving Ax=b. In this talk, we prove that MINRES (and CG on the normal

equations) achieves an optimal complexity (up to a factor of 4) over all randomized matrix-vector methods, and GMRES achieves the best possible over all "transpose-free" methods.

Polynomials, Twin BiCG and Approximating the Inverse

Author: Ron Morgan[1]

[1] *Baylor University*

Contact: ronald_morgan@baylor.edu

The BiCG method for solving linear equations has a polynomial at its core. The new Twin BiCG method solves multiple right-hand systems using the same polynomial for each system. This polynomial is applied implicitly by using the parameters from solving the first right-hand side for all of the systems. Twin BiCG has automatic stability control from the extra copies of eigenvalues that are produced by the nonsymmetric Lanczos algorithm.

We also discuss how Twin BiCG can give an approximation to the inverse of a matrix. Applications include Multilevel Monte Carlo for determining the trace of the inverse of a large matrix. Also, shift-and-invert Arnoldi for computing eigenvalues can solve the associated linear equations to less accuracy as the iteration proceeds. Twin BiCG can efficiently implement this.

On block Krylov and matrix polynomials

Authors: Michele Rinelli[1], Raf Vandebril[2]

[1] *KU Leuven*

[2] *Dept. Computer Science, KU Leuven*

Contact: michele.rinelli@kuleuven.be

The deep connection between Krylov methods, scalar orthogonal polynomials, and moment matrices is well established, particularly for Hermitian and unitary matrices. In this talk, we extend this framework to block Krylov methods and orthogonal matrix polynomials.

By representing the elements of a block Krylov subspace via matrix polynomials, we consider the matrix-valued inner product introduced in [1], which, under a non-degeneracy assumption, defines a linear isometry. This establishes a one-to-one correspondence between orthonormal matrix polynomials and orthonormal bases of the block Krylov subspace.

For normal matrices, the block Gauss discretization [1], [2] of such an inner product admits an integral representation familiar from the theory of orthogonal matrix polynomials. As an application, we extend a Szegő-type short recurrence, originally developed for matrix polynomials on the unit circle [3], to the block Arnoldi algorithm applied to unitary matrices.

Finally, we analyze the structure of the block moment matrix and explore its connection to orthogonal matrix polynomials and recurrence coefficients via a Cholesky-like factorization.

[1] Lund, K.: A New Block Krylov Subspace Framework with Applications to Functions of Matrices Acting on Multiple Vectors, Phd thesis, Temple University and Bergische Universität Wuppertal, (2018)

[2] Zimmerling, J., Druskin, V., Simoncini, V.: Monotonicity, bounds and acceleration of block Gauss and Gauss-Radau quadrature for computing $B^T\phi(A)B$, J. Sci. Comput., 103 (2025)

[3] Sinap, A., Van Assche, W.: Orthogonal matrix polynomials and applications, J. Comput. Appl. Math. 66 (1996)

(Block) Lanczos Function Approximation for Quasi-Newton Optimization Algorithms

Author: Cooper Simpson[1]

[1] *University of Washington*

Contact: rscooper@uw.edu

The Lanczos process is a well-known Krylov subspace method for the orthogonal tridiagonalization of a hermitian matrix $\mathbf{Z}$. Equipped with a suitable function f, Lanczos function approximation (LFA) can be used as a powerful tool for approximating the matrix-function $f(\mathbf{Z})$ or matrix-function-vector products $f(\mathbf{Z})\omega$.

We discuss an application of LFA in computing iterates of the form

$$\mathbf{x}_{(k+1)} = \mathbf{x}_{(k)} - \eta_{(k)} \left(\left(\nabla^2 g(\mathbf{x}_{(k)}) \right)^2 + \lambda_{(k)} \mathbf{I} \right)^{-1/2} \nabla g(\mathbf{x}_{(k)})$$

arising in the regularized saddle-free Newton (R-SFN) algorithm for $g \in \mathcal{C}^2$. LFA facilitates efficient update steps yielding competitive performance against adaptive regularization with cubics (ARC), which is among the best available algorithms for non-convex optimization.

Additionally, we will detail ongoing work for improving LFA in the general setting, motivated by our application in optimization. We will discuss residual estimation for iterative refinement and block Lanczos for implicit preconditioning.

Preconditioned log-determinant approximation: one probe vector is almost always enough!

Authors: Alice Cortinovis[1], Daniele Toni[2]

[1] *University of Pisa*

[2] *Scuola Normale Superiore*

Contact: daniele.toni@sns.it

We present randomized algorithms for estimating the log-determinant of regularized symmetric positive semi-definite matrices. The algorithms access the matrix only through matrix vector products, and are based on the introduction of a preconditioner and stochastic trace estimator. We claim that preconditioning as much as we can and making a rough estimate of the residual part with a small budget achieves a small error in most of the cases. We choose a Nyström preconditioner and estimate the residual using only one sample of stochastic Lanczos quadrature. We analyze the performance of this strategy from a theoretical and practical viewpoint. We also present an algorithm that, at almost no additional cost, detects whether the proposed strategy is not the most effective, in which case it uses more samples for the stochastic Lanczos quadrature part. Numerical examples on several test matrices show that our proposed methods are competitive with existing algorithms.

A Generalized Framework for Orthogonal Rational Functions applied to Rational Approximation

Author: Robbe Vermeiren[1]

[1] *KU Leuven*

Contact: robbe.vermeiren@kuleuven.be

Rational functions are fundamental to several non-linear approximation problems in, for example, model reduction, system identification, and PDE problems. Consequently, one is often interested in constructing an orthonormal basis of rational functions to ensure numerical stability and conditioning.

In this talk, we present a generalized framework for constructing such bases for rational function *vectors* of arbitrary length k. We show that the underlying recurrence relations can be represented by a specific structured pair of k-Hessenberg matrices. We propose two efficient algorithms to compute these bases:

1. An updating algorithm using rotations,
2. A rational Arnoldi-type algorithm.

To demonstrate the robustness of this framework, we apply it to the rational approximation of $\sqrt{z}$ on $[0, 1]$. We show that our method successfully handles exponentially clustered poles and recovers the optimal lightning + polynomial convergence rates.

Everything is Vecchia: Unifying low-rank and sparse inverse approximations

Authors: Eagan Kaminetz[1], Robert Webber[1]

[1] *University of California San Diego*

Contact: rwebber@ucsd.edu

The partial pivoted Cholesky approximation accurately represents matrices that are close to being low-rank. Meanwhile, the Vecchia approximation accurately represents matrices with inverse Cholesky factors that are close to being sparse. What happens if a partial Cholesky approximation is combined with a Vecchia approximation of the residual? We show how the sum can be rewritten as a Vecchia approximation of the original matrix with an amplified sparsity pattern. Thus, the Vecchia approximation is a superset for other factored matrix approximations and it has broad applicability.

MS 37: New Advancements in Tensor Decomposition and Computation

Minisymposium Organizers: Anna Konstorum, Carmeliza Navasca

In this mini symposium, we gather experts in tensor decomposition and multilinear algebra who are tackling challenging problems in data science, machine learning, image processing, dynamical systems, control systems and the biomedical sciences. We will discuss topics in multilinear dynamical system theory, model order reduction, tensor completion, symmetric tensor decomposition, algorithms for PDEs and quantum simulations.

List of Talks

Preconditioners for Kronecker Sum Systems with Applications to Polynomial Feedback Control

Authors: Jeff Borggaard[1], Hamza Adjerid[1], Ali Bouland[1]

[1] *Virginia Tech*

Contact: jborggaard@vt.edu

Multivariate polynomial approximations to Hamilton-Jacobi-Bellman equations can be expressed using Kronecker products leading to very large, but structured, linear systems. Their structure appears as n-way generalizations of Lyapunov or generalized Lyapunov equations. For monomial terms of degree d, their dimension scales as the number of state dimensions n raised to the d. For feedback control problems with modest dimension n=*1000* and relatively low degree approximations of d=*3* or *4*, matrix-free iterative solvers using scalable preconditioners are essential. This talk

will review several options for preconditioners, then present a number of examples arising from discretizations of nonlinear partial differential equation control problems that demonstrate the performance of these linear solvers. As expected, there can be a strong connection between the discretization and the preconditioner choices. We conclude with applications to fluid flow control problems modeled through approximations to Navier-Stokes equations.

A greedy approach for approximate tensor diagonalization

Authors: Anna Konstorum, Anna Ma[1]

[1] *University of California, Irvine*

Contact: konstorum.anna@gmail.com

Symmetric tensor diagonalization has applications in statistics and signal processing. Unlike for real symmetric matrices, there is no guarantee that a real-valued symmetric tensor is diagonalizable. Therefore, one generally approaches the problem as an approximate tensor diagonalization (ATD) problem. In this talk, we show that Jacobi-type methods for ATD that naturally extend the Jacobi method for real symmetric matrix diagonalization share several useful properties with the original method, including the conservation of mass during the iteration steps. We use these properties to generalize the greedy method for the matrix case to third-order real-valued tensors, and present numerical results comparing our approach to the classical sweep method. We discuss convergence properties and place our work in the context of Jacobi-type approaches for ATD.

The Nondecreasing Rank

Author: Andrew McCormack[1]

[1] *University of Alberta*

Contact: mccorma2@ualberta.ca

Matrix or tensor data often has structured rows, columns, or more generally modes. In particular, a mode may have a natural ordering that can be leveraged to obtain parsimonious representations of the data. To this end, the concept of the nondecreasing (ND) rank is introduced in this talk. A tensor has an ND rank of r if it can be represented as a sum of r outer products of vectors, with each vector satisfying a monotonicity constraint. It is shown that for certain poset orderings finding an ND factorization of rank r is equivalent to finding a nonnegative rank r factorization of a transformed tensor. However, not every tensor that is monotonic has a finite ND rank. Theory is developed describing the properties of the ND rank, including typical, maximum, and border ND ranks. Highlighted also are the special settings where a matrix or tensor has an ND rank of

one or two. As a means of finding low ND rank approximations to a data tensor a variant of the hierarchical alternating least squares algorithm is introduced.

Transform-Based Multilinear Algebra via Tensor Decompositions

Authors: Can Chen[1], Yidan Mei[2]

[1] *University of North Carolina at Chapel Hill*

[2] *Yale University*

Contact: yidan.mei@yale.edu

Transform-based tensor products, including the T-product and its more general form, namely the higher-order tensor-tensor product, have been widely used in image processing, signal reconstruction, and robotics. While invertible transforms enable tensor computations to be carried out via matrix operations in the transform domain, the resulting storage and computational costs remain prohibitive for high-dimensional tensors. To address this challenge, we leverage tensor decomposition techniques, including tensor train decomposition (TTD) and hierarchical Tucker decomposition (HTD), to accelerate transform-based multilinear algebra. In particular, we develop TTD- and HTD-based formulations for the T-product and its associated algebra by operating directly on the factor matrices or tensors of the decompositions. The framework is further generalized to the higher-order tensor-tensor product and applied to multilinear control problems. We demonstrate the effectiveness of our framework with numerical examples.

Tensor Data for Control Strategies in Systems

Authors: Carmeliza Navasca[1], Jiahua Jiang[2]

[1] *University of Alabama at Birmingham*

[2] *University of Birmingham*

Contact: cnavasca@uab.edu

We study optimal control problems arising from partial differential equations. More specifically, we look at the optimal control of Allen-Cahn Equation (ACE) with a source term. ACE is well known for modeling phase transitions and thus, has many applications, apart from physics (semiconductors), in biological systems, material science and image processing. ACE models cellular membranes which are vital in transporting vesicles, viral budding and dynamical sorting and signaling.

Through the dynamic programming principle, optimal control problem are reformulated into the Hamilton-Jacobi-Bellman (HJB) Equations. It is well known that the computational cost is prohibitively expensive when standard methods are implemented in solving the HJB Equations. Thus we obtain reduced order models, using tensor decomposition techniques for reduced order basis.

Sparse optimization and flexible hybrid methods are used for low rank canonical polyadic tensor decomposition. The reduced optimal control problem leads to reduced state-dependent Riccati Equations which can be solved efficiently.

Revisit CP Tensor Decomposition: Statistical Optimality and Fast Convergence

Author: Augustine (Runshi) Tang[1]

[1] *University of Wisconsin-Madison*

Contact: rtang56@wisc.edu

Canonical Polyadic (CP) tensor decomposition is a fundamental technique for analyzing high-dimensional tensor data. While the Alternating Least Squares (ALS) algorithm is widely used for computing CP decomposition due to its simplicity and empirical success, its theoretical foundation, particularly regarding statistical optimality and convergence behavior, remain underdeveloped, especially in noisy, non-orthogonal, and higher-rank settings.

In this work, we revisit CP tensor decomposition from a statistical perspective and provide a comprehensive theoretical analysis of ALS under a signal-plus-noise model. We establish non-asymptotic, minimax-optimal error bounds for tensors of general order, dimensions, and rank, assuming suitable initialization. To enable such initialization, we propose Tucker-based Approximation with Simultaneous Diagonalization (TASD), a robust method that improves stability and accuracy in noisy regimes. Combined with ALS, TASD yields a statistically consistent estimator. We further analyze the convergence dynamics of ALS, identifying a two-phase pattern—initial quadratic convergence followed by linear refinement. We further show that in the rank-one setting, ALS with an appropriately chosen initialization attains optimal error within just one or two iterations.

Tensor Methods for Multi-omics Data

Author: Neriman Tokcan[1]

[1] *University of Massachusetts Boston*

Contact: neriman.tokcan@umb.edu

High-throughput genomics and omics technologies generate data with intrinsic multi-way structure arising from multiple samples, molecular features, experimental conditions, and biological contexts. Standard matrix-based methods often obscure this structure through flattening or aggregation. Tensor-based representations provide a natural mathematical framework for preserving and exploiting the underlying multilinear organization of such data.

In this talk, we present tensor decomposition models for the analysis of high-dimensional omics datasets and demonstrate how these models yield low-rank representations that enable dimensionality reduction while maintaining interpretability through mode-specific latent factors. Applications include multi-sample gene expression analysis, multi-omics integration, and higher-order relational modeling of cell–cell interactions using multi-omics data.

We discuss key mathematical and computational challenges arising in these applications, including rank determination, identifiability and uniqueness, robustness to noise, scalability, and the incorporation of structural constraints motivated by domain knowledge. We conclude by outlining open problems at the interface of multilinear algebra, optimization, and statistical modeling, highlighting opportunities for theoretical advances driven by modern omics data.

Multiscale Grassmann Manifolds for Single-Cell Data Analysis

Author: Xiang Xiang Wang[1]

[1] *Michigan State University*

Contact: wangx306@msu.edu

Single-cell data are typically represented in Euclidean space, which limits the ability to capture intrinsic correlations and multiscale geometric structure. We propose a multiscale framework based on Grassmann manifolds that represents single-cell RNA-seq data through subspace geometry across multiple scales.

MS 38: Combinatorial Matrix Theory

Minisymposium Organizers: Lei Cao, Louis Deaett

Combinatorial matrix theory studies families of matrices defined by some combinatorial conditions. These can be studied as combinatorial objects in their own right, as with -matrices, alternating sign matrices, etc., while they also can provide a combinatorial description of general matrices over rings and fields, as in the context of sign patterns, Markov chains, etc., where the relationship between the combinatorics of the matrices and properties of the operators they represent is often the focus, for example in considering what the sign pattern of a square matrix implies about its eigenvalues. In all cases, connections with other combinatorial objects, such as simple/directed/bipartite graphs, permutations, matroids, etc., are central. The goal of this session is to feature work representing a spectrum of investigation in combinatorial matrix theory and its applications to other areas.

List of Talks

Hermitian Spectra of Merged Subdivision Mixed Graphs

Authors: Eber Lenes[1], Enide Andrade[2], Geraldine Infante[3], María Robbiano[3]

[1] *Área de Ciencias Básicas Exactas, Universidad del Sinú, Cartagena, Colombia*

[2] *Center for Research and Development in Mathematics and Applications (CIDMA), Department of Mathematics, University of Aveiro, Portugal*

[3] *Department of Mathematics, Universidad Católica del Norte, Chile*

Contact: enide@ua.pt

Let G be a mixed graph and let (H_1, H_2) be an ordered pair of mixed graphs whose orders coincide with the order and size of G, respectively. We introduce the subdivision mixed graph $S(G)$ and the (H_1, H_2)-merged subdivision mixed graph. We investigate the Hermitian spectrum and the Hermitian energy of these graphs, deriving spectral properties that relate merged subdivision mixed graphs to their underlying mixed and undirected counterparts. In particular, we compare the Hermitian energies of certain classes of mixed graphs with those of their underlying graphs, illustrating how the merging process influences spectral behavior. These results show how graph operations influence Hermitian spectral invariants in mixed graphs.

This is joint work with G. Infante, M. Robbiano, Universidad Católica del Norte, Chile,and Eber Lenes, Universidad del Sinú, Colombia

---

A new combinatorial rank concept and its challenges

Authors: Geir Dahl, Richard Brualdi[1]

[1] *University of Wisconsin*

Contact: geird@math.uio.no

We introduce a new rank concept for $(0, \pm 1)$-matrices, called the $\pm$-rank of a $(0, \pm 1)$-matrix. This "generalizes" the binary rank and the term rank of (0,1)-matrices. We establish several inequalities relating the different ranks, including ordinary real rank. Moreover, the $\pm$-rank is discussed for certain classes of $(0, \pm 1)$-matrices, such as alternating sign matrices (ASMs), network matrices etc. We also focus on some challenging problems concerning this rank concept, both of theoretical and computationally nature.

Geometric and Combinatorial Properties of the ASM Polytope

Authors: Elizabeth Dinkelman[1], Walter Morris[1]

[1] *George Mason University*

Contact: elizabethdinkelman@hotmail.com

The polytope ASM_n, the convex hull of the $n \times n$ alternating sign matrices, was introduced by Striker and by Behrend and Knight. A face of ASM_n corresponds to an elementary flow grid defined by Striker, and each elementary flow grid determines a doubly directed graph defined by Brualdi and Dahl. We show that a face of ASM_n is symmetric if and only if its doubly directed graph has all vertices of even degree. We show that every face of ASM_n is a 2-level polytope. We show that a d-dimensional face of ASM_n has at most 2^d vertices and $4(d-1)$ facets, for $d \geq 2$. We show that a d-dimensional face of ASM_n satisfies $vf \leq d2^{d+1}$, where v and f are the numbers of

vertices and facets of the face. If the doubly directed graph of a d-dimensional face is 2-connected, then $v \leq 2^{d-1} + 2$. We describe the facets of a face and a basis for the subspace parallel to a face in terms of the elementary flow grid of the face. We prove that no face of ASM_n has the combinatorial type of the Birkhoff polytope B_3. We list the combinatorial types of faces of ASM_n that have dimension 4 or less.

The sum of a topological index and its reciprocal index

Author: Wei Gao
Contact: wvg5121@psu.edu

Let G be a simple connected graph. A vertex-degree-based topological index is defined as

$$TI_f(G) = \sum_{uv \in E(G)} f(d_u, d_v),$$

where $f(x, y)$ is a symmetric real function. In theoretical chemistry, these indices serve as essential numerical molecular descriptors in QSAR/QSPR models. In this work, we investigate the extremal properties of $TI_f + RTI_f$, defined as the sum of a topological index and its reciprocal. Focusing on the first Zagreb index ($f(x, y) = x + y$), the second Zagreb index ($f(x, y) = xy$), and the forgotten index ($f(x, y) = x^2 + y^2$), we characterize the graphs that achieve the maximum and minimum values of $TI_f + RTI_f$ among all trees. Furthermore, we extend our analysis to the extremal problem for $TI_f + RTI_f$ of k-uniform hypergraphs.

Frobenius theorem for Cullis determinant

Authors: Alexander Guterman[1], Andrey Yurkov[1]

[1] *Bar-Ilan University*

Contact: alexander.guterman@gmail.com

The first results on transformations preserving matrix invariants is due to Frobenius. This result describes the structure of linear maps T preserving the determinant function, i.e., $\det X = \det T(X)$ for all X. Later on there were several extension of this result which are due to Diedonnie, Schur, Dynkin and others.

In 1913 Cullis and then in 1966 independently Radić introduced an analog of the determinant function for the rectangular matrices. Namely, given an injection $\sigma \in S_K^I$, we denote by $\mathrm{sgn}_{nk}(\sigma)$ the product

$$s \cdot (-1)^{\sum_{l=1}^{k} (\sigma(l) - l)},$$

where $s = \operatorname{sgn}(\pi)$ is the sign of the permutation

$$\pi = \begin{pmatrix} i_1 & \dots & i_k \\ \sigma(1) & \dots & \sigma(k) \end{pmatrix},$$

here $\sigma(K) = \{i_1, \dots, i_k\}$, and $i_1 < i_2 < \dots < i_k$. Let $n \geq k$, $X \in M_{nk}(\mathbb{F})$. Then Cullis determinant $\det_{n\,k}(X)$ of X is defined to be the function:

$$\det_{n\,k}(X) = \sum_{\sigma \in S_k^I} \operatorname{sgn}_{nk}(\sigma) X_{(\sigma(1),1)} X_{(\sigma(2),2)} \dots X_{(\sigma(k),k)}.$$

Cullis determinant has many interesting properties of the usual determinant as well as many new properties. It is interesyting as from theoretical point of view, as due to the number of applications. In this talk we discuss an analog of Frobenius theorem for Cullis determinant, namely we characterize linear preservers of this function. It appears that the preservers depend on the parity of $k + n$. Our proof is based on analogs of Flander's results and matroid theory.

The talk is based on the series of joint works with Andrey Yurkov.

The Rank-Preserving Transversality Property

Authors: Marina Arav[1], Frank Hall[1], Chris Hart[1], Hein van der Holst[1], Zhongshan Li[1], Zixuan Li[1], Jian Liu[1], Jiamin Pan[1], Joonwon Seo[1], Li Wang[1], Hanfei Xu[1], Yiran Xu[1], Zheng Yang[1], Yue Zhao[1]

[1] *Georgia State University*

Contact: ahart13@gsu.edu

Let A be an $m \times n$ real matrix. If the manifolds $\widetilde{\mathcal{M}}_A = \{H^{-1}AG : G, H \text{ are nonsingular}\}$ and $Q(\operatorname{sgn}(A))$ intersect transversally at A, that is, the tangent spaces of $\widetilde{\mathcal{M}}_A$ and $Q(\operatorname{sgn}(A))$ at A sum to $\mathbb{R}^{m\times n}$, we say that A has the rank-preserving transversality property (RPTP) and that A is an RPTP matrix. We establish many important properties of RPTP matrices as well as several sign pattern classes that require the RPTP. For example, RPTP matrices are closed under permutation equivalence, diagonal equivalence, and transpose. Further, a block upper triangular matrix with all diagonal blocks square has the RPTP if and only if each diagonal block has the RPTP and at most one diagonal block is singular. Just as the Strong Spectral Property is useful in studying the spectra of symmetric matrices associated with a graph, the notion of RPTP is a useful tool for studying the minimum ranks of sign patterns and zero-nonzero patterns.

How to Combinatorially Prove Matrices are Mutually Inverse

Authors: Aditya Khanna[1], Nicholas Loehr[1]

[1] *Virginia Tech*

Contact: adityakhanna@vt.edu

Let $A = (A_n)_{n\geq 0}$ and $B = (B_n)_{n\geq 0}$ be families of "recursive" rectangular matrices, that is, the entries of A_n can be expressed as a linear combination of entries of A_m for $m < n$, and similarly for B. Such families arise in algebraic combinatorics as change-of-basis matrices between the basis of symmetric functions and their generalizations. The entries of such matrices can be computed as sums over signed-weighted combinatorial objects which themselves can be constructed recursively. In this talk, I will describe a framework that reduces the condition $A_nB_n = I$ for all $n \geq 0$ on the matrix families to an equivalent local condition on the objects which can be handled using combinatorial methods such as sign-reversing involutions. This framework has applications in producing canonical bijective proofs for the Kostka matrix case, the orthogonality of symmetric group characters, and Mobius inversion for the lattice of sets. All combinatorial concepts will be defined in the talk.

Spectral bounds for the clique number of a graph

Authors: Sudipta Mallik[1], Hareshkumar Jadav[2], Priyanshu Pant[2], Ranveer Singh[2]

[1] *Marshall University*

[2] *IIT Indore, India*

Contact: mallik@marshall.edu

Given an integer k, deciding whether a graph has a clique of size k is an NP-complete problem. Wilf's inequality provides a spectral lower bound for the clique number (i.e., the order of a largest clique) in terms of the largest adjacency eigenvalue. In 2018, Elphick and Wocjan conjectured a stronger spectral bound using positive adjacency eigenvalues. We introduce a spectral bound using Laplacian eigenvalues. We show that this is an improvement for almost all graphs.

On the Length of a Multiplicity List of a Graph

Author: Johnna Parenteau
Contact: parenjoh@uregina.ca

Given a simple graph, G, on n vertices, let $S(G)$ be the set of $n \times n$ real symmetric matrices, $A = [a_{ij}]$, associated to G where, when $i \neq j$, $a_{ij} \neq 0$ if and only if ij is an edge in G and the main diagonal is free to be chosen. For any square matrix, A, let $q(A)$ equal the number of distinct eigenvalues of A. The minimum number of distinct eigenvalues of G is defined as $q(G) = \min\{q(A) : A \in S(G)\}$. A natural question that arises from this association is: Does there exist an $A \in S(G)$ with $q(A) = k$ for every integer k in the interval $[q(G), n]$? Since the inverse eigenvalue problem for graphs has been studied from many angles for years, a general solution seems

difficult, so any progress made on subproblems, such as $q(G)$, contributes to an understanding of the whole problem. In this talk, we discuss some recent progress on this question for certain classes of graphs.

Sparsity of a matrix and its inverse

Author: Bryan Shader[1]

[1] *University of Wyoming*

Contact: bshader@uwyo.edu

This talk discusses problem of determining the minumum number of nonzero entries in a pair of matrices (A,A^(-1)) for A in various families of matrices (e.g. irreducible, fully indecomposable, primitive, positive definite, orthogonal, symmetric with connected graph, irreducible covariance matrices). Some of the results are from work with H. Gupta, L. Hogben and T. Wong.

Combinatorial characterization of matrix algebras over finite fields

Authors: Nik Stopar[1], Alen Đurić[2], Sara Koljančić[2]

[1] *University of Ljubljana*

[2] *University of Banja Luka*

Contact: nik.stopar@fgg.uni-lj.si

In this talk we demonstrate how a matrix algebra over a finite field can be completely described using combinatorial properties. The main tool that allows one to do this is the compressed zero-divisor graph of a ring, which describes pairs of matrices A and B such that $AB = 0$. We list a set of 5 combinatorial axioms that uniquely determine the compressed zero-divisor graph $\Theta(M_n(\mathbb{F}))$ of the matrix algebra $M_n(\mathbb{F})$, where $|\mathbb{F}| = p^m$. Furthermore, the structure of this graph uniquely determines the ring $M_n(\mathbb{F})$ itself up to isomorphism. We also discuss some properties of the compressed zero-divisor graphs that may be useful for investigating subalgebras of $M_n(\mathbb{F})$ and individual matrices.

MS 39: Application-driven Family of Matrix Computations: Factorization, Inverse, Linear Solve

Minisymposium Organizer: Kapil Ahuja

The real-life application, where a matrix arises, plays a big role in efficiently performing the underlying matrix computation. Some of the leading applications of the current times are fracture mechanics, language modelling, and high performance computing. This minisymposium explores recent advances in family of matrix computations involving factorization, inverse, and linear Solves for these applications.

List of Talks

Scalable Approximate Selected Inversion via ILU and Spectral Corrections for Sparse Systems

Authors: Tahamina Akter[1], Matthias Bollhoefer[1]

[1] *TU Braunschweig*

Contact: tahamina.akter@tu-braunschweig.de

We analyze two parallel numerical strategies for computing selected entries of the matrix inverse of large, sparse, symmetric systems: The selected inverse method and a factorized approximate inverse method. Both techniques are aimed at computations via LU factorizations or incomplete LU (ILU) factorizations. The selected inverse approach exploits the LU/ILU factorization to recover the inverse within the pattern of the (incomplete) LU factorization. In contrast, the factorized approximate inverse method applies a truncated series expansion (Neumann-type), providing an alternative approach at the cost of reduced accuracy. To improve accuracy while keeping sparsity,

we introduce low-rank corrections and an adaptive eigenvector deflation strategy. This hybrid approach sets tightening drop tolerances off when possible and helps balance accuracy, fill-in, and computation time. Numerical illustrations for both parallel and sequential computations demonstrate the effectiveness and robustness of our approach.

Loracx: Low Rank Approximations with Constraints at Exascale

Author: Ramakrishnan Kannan
Contact: kannanr@ornl.gov

Analyzing large-scale scientific data—such as molecular dynamics simulations of MoS_2 recrystallization—poses significant challenges for traditional methods like Nonnegative Matrix Factorization (NMF), particularly on exascale systems. In this talk, we introduce Low-Rank Approximations with Constraints at Exascale (LORACX), a scalable framework that employs distributed, GPU-accelerated NMF integrated into a modern, Python-based HPC stack.

Key innovations include communication-efficient designs using blocked and overlapped algorithms to mitigate latency and memory constraints, as well as GPU-optimized Nonnegative Least Squares (NNLS) solvers. Performance evaluations on up to 8,192 Frontier nodes demonstrate strong scalability, processing a 16.3 $\times$ 16.3 million matrix in 3 seconds and achieving 0.67 exaflops in double precision. We present detailed weak-scaling results, including computational versus communication cost analyses, and show that baseline comparisons consistently confirm the superior performance of LORACX-GPU.

Applied to MoS_2 molecular dynamics data, LORACX successfully identifies structural motifs and captures phase transition dynamics, highlighting its potential as a powerful tool for large-scale materials science discovery.

Exploiting Kronecker Structure and Krylov Subspaces for Scalable Second-Order Optimization in Federated Physics-Informed Neural Networks

Authors: Amit Upadhyay[1], Dharavath Ramesh[1]

[1] *Indian Institute Of Technology (Indian School Of Mines) Dhanbad*

Contact: 23dr0016@iitism.ac.in

Solving partial differential equations (PDEs) using distributed Physics-Informed Neural Networks (PINNs) introduces major computational challenges associated with high-dimensional curvature estimation, ill-conditioned optimization landscapes, and communication overhead in federated environments. In this work, we exploit Kronecker structure and Krylov subspace methods to develop

a scalable second-order optimization framework for federated PINNs applied to the 2D Poisson equation. The proposed method combines Kronecker-Factored Approximate Curvature (K-FAC) with Conjugate Gradient (CG) iterations within a matrix-free Gauss–Newton formulation, avoiding explicit Hessian or Jacobian construction through automatic differentiation-based Jacobian-vector products. The symmetric positive definite structure induced by discretized Poisson operators supports efficient CG convergence while enabling scalable curvature-aware optimization. Within the federated learning framework, each client independently solves a local Gauss–Newton system using K-FAC approximations and Krylov iterations, transmitting only local parameter corrections to the central server. Experimental results across grid resolutions from 32×32 to 128×128 demonstrate up to 57.8% relative improvement in test accuracy and 44.4% reduction in required communication rounds compared with first-order Adam-based optimization. These findings illustrate how structured linear algebra techniques, particularly Kronecker factorization and Krylov subspace solvers, provide an efficient foundation for communication-efficient scientific machine learning under distributed computing constraints.

Matrix-Free Geometric Multigrid Preconditioning Of Combined Newton-GMRES For Solving Phase-Field Fracture With Local Mesh Refinement

Authors: Leon Kolditz[1], Sanghyun Lee[2], Thomas Wick[1]

[1] *Leibniz University Hannover*

[2] *Florida State University*

Contact: thomas.wick@ifam.uni-hannover.de

In this talk, the matrix-free solution of quasi-static phase-field fracture problems is further investigated. More specifically, we consider a quasi-monolithic formulation in which the irreversibility constraint is imposed with a primal-dual active set method. The resulting nonlinear problem is solved with a line-search assisted Newton method. Therein, the arising linear equation systems are solved with a generalized minimal residual method (GMRES), which is preconditioned with a matrix-free geometric multigrid method including geometric local mesh refinement. Our solver is substantiated with a numerical test on locally refined meshes. Afterwards, we discuss applications of this solver framework in fluid-filled fractures in porous media.

Contributed Talks

List of Talks

Quantum Spectral Methods for Poisson and Heat Equations

Authors: Chih-Kang Huang[1], Giacomo Antonioli[2]

[1] *Inria, Côte d'Azur University*

[2] *Universita di Pisa*

Contact: giacomo.antonioli@phd.unipi.it

We present a quantum framework for solving a large class of elliptic and parabolic Partial Differential Equations (PDEs) endowed with periodic conditions. We solve the Poisson equation $\Delta u = f$ and the heat equation on the d-dimensional flat torus using a Fourier spectral method implemented on quantum circuits.

The main contribution is an efficient use of block encoding to load the inverse diagonal spectral filter into a quantum circuit. Unlike standard quantum algorithms based on general matrix inversion, our approach constructs the reciprocal of the diagonal entries directly. This procedure requires only one ancillary qubit and can be implemented in polylogarithmic time with respect to the system size N.

Different strategies are adopted for solving the problem, depending on the type of equation. For elliptic equations, the solution is obtained by directly applying the inverse spectral filter to the source term. For parabolic equations, we employ an iterative implicit scheme where the spatial variables are handled via our quantum spectral method, while the update for the next time step is performed classically. We validate our approach by comparing the energy evolution and steady-state solutions against standard classical spectral schemes. The numerical results demonstrate that our quantum solver achieves high fidelity, with errors comparable to machine precision limits.

Pretty good state transfer in Grover walks on abelian Cayley graphs

Authors: Koushik Bhakta[1], Bikash Bhattacharjya[1]

[1] *Indian Institute of Technology Guwahati*

Contact: b.koushik@iitg.ac.in

In this work, we study pretty good state transfer (PGST) in Grover walks on graphs. We consider the transfer of quantum states localized at the vertices of a graph and use Chebyshev polynomials to analyze PGST between such states. In general, we find a necessary and sufficient condition for the occurrence of PGST on graphs. We then focus our analysis on abelian Cayley graphs and derive a necessary and sufficient condition for the occurrence of PGST on such graphs. Consequently, we

obtain a complete characterization of PGST on unitary Cayley graphs. Our results yield infinite families of graphs that exhibit PGST yet do not exhibit perfect state transfer.

Spectral clustering and accelerated SSO for large and dense problem

Authors: Elias Jarlebring[1], Jiayu Bian[2]

[1]*KTH*

[2]*KTH Royal Institute of Technology*

Contact: jiayub@kth.se

Spectral clustering is a graph-based method for data partitioning, which relates a relaxed graph partitioning problem to the eigenvectors of the associated graph Laplacian; see, e.g., [Peng et al., SIAM J. Comput. 46(2):710–743, 2017]. Following the computation of these eigenvectors, a post-processing step is used to determine the final clustering. Typically, k-means is the standard choice of post-processing algorithm. Recently [Eldén, SIAM J. Matrix Anal. Appl. 45(1):112–133, 2024], an algorithm called Semi-sparse Orthogonal (SSO) Approximation, has been proposed as an alternative that improves on both clustering accuracy and stability. Given the eigenvector matrix $X \in R^{n\times p}$, SSO seeks $\min_{Q,\Psi} \|XQ - \Psi\|_F$, where $Q \in R^{p\times p}$ is orthogonal and $\Psi \in R^{n\times p}$ is an indicator matrix with one nonzero per row. The method alternates between updating Ψ by assigning each row to the entry of the largest magnitude of XQ, and updating Q by the polar factor of $\Psi^T X$.

In this work, we consider clustering problems on large and dense graphs, with a large number of clusters. In such settings, standard Arnoldi type solvers become computationally unattractive. Furthermore, computing the polar factor based on singular value decomposition becomes infeasible at these scales. We present an ongoing investigation that circumvents these difficulties by replacing explicit eigenvector computations with the spectral projector of the graph Laplacian, and by performing post-processing via a large-scale variant of SSO. Our approach exploits methods from matrix functions to compute the spectral projectors and the polar factors using polynomial approximations. This avoids the explicit computation of individual eigenvectors and is well suited for large-scale, dense problems.

Linear Geometry Insights in the Expectation Maximization Algorithm Convergence

Authors: Dacian Bonta[1], John Aarsvold[1]

[1]*Emory University*

Contact: dacian.bonta@va.gov

The expectation maximization (EM) algorithm, widely used in medical image tomographic reconstruction, iteratively estimates the values of a discrete distribution of emission sources $\rho_{t=1,2,..,T}$ from a set of detector measurements $A_{k=1,2,..,K}$. The two steps of an iteration are:

$$A_k^{(n)} = \sum_t P_{k\leftarrow t} \cdot \rho_t^{(n)} \text{ (E) and } \rho_t^{(n+1)} = \rho_t^{(n)} \sum_k \left[\frac{A_k}{A_k^{(n)}} P_{k\leftarrow t} \right] \text{ (M)}$$

where superscript indicates estimated values at step n, and $P_{k\leftarrow t}$ is the probability that an emission from source t is detected by detector k.

Using analytical tools, it has been shown that the algorithm progresses to a fixed point, but further description of the ideal algorithm behavior is incomplete.

We propose to view measurements A_k as defining hyperplanes in the finite-dimensional linear space defined by the emission source values $(\rho_{t=1,2,..,T})$ (equation E). In imaging the system is overdetermined (more measurements than emission sources), thus, when the attached linear system is consistent, these hyperplanes have a single common point (the solution).

Because all values appearing in equations (E) and (M) are positive, we argue adjustments to ρ_t in step (M) induced by measurements k with $\frac{A_k}{A_k^{(n)}} > 1$ cannot cancel adjustments induced by measurements k with $\frac{A_k}{A_k^{(n)}} < 1$. We show the two convex cones defined by the direction vectors to the $\frac{A_k}{A_k^{(n)}} > 1$ hyperplanes and the direction vectors to the $\frac{A_k}{A_k^{(n)}} < 1$ hyperplanes, respectively, with the vertex in the solution point, have only that point as intersection. Thus, the corresponding adjustments cannot cancel each other and, for a consistent overdetermined problem, the only fixed point for the EM algorithm is the solution.

The linear geometry view of the EM algorithm promises a path for better understanding of this important algorithm.

A More Robust Streaming Dynamic Mode Decomposition

Authors: Ela Đimoti[1], Zlatko Drmac[2]

[1] *University of Zagreb*

[2] *Faculty of Science, University of Zagreb*

Contact: ela.dimoti@math.hr

Dynamic Mode Decomposition (DMD) is a data-driven tool for capturing complex nonlinear dynamics. It can be used to identify, analyze and forecast dynamical systems $x_{k+1} = F(x_k)$ governed by an unknown or complex mapping F using only observed snapshots $s_1, s_2, ..., s_{n+1}$. If we denote $X := (s_1\ s_2\ \cdots\ s_n)$, $Y := (s_2\ s_3\ \cdots\ s_{n+1})$, finite-dimensional approximation of the governing function F is a matrix A such that $AX \approx Y$ in the least-squares sense.

In a streaming application DMD has to be recomputed over a widening data window. If we denote a rank-revealing orthonormal decomposition of $X = Q_x\tilde{X}$ and $Y = Q_y\tilde{Y}$, Hemati et al. (2014)

propose updating Rayleigh-Ritz matrix

$$Q_x^T A Q_x = (Q_x^T Q_y)(\tilde{Y}\tilde{X}^T)(\tilde{X}\tilde{X}^T)^{-1} =: (Q_x^T Q_y) G_{y,x} G_x^{-1}$$

via updates of $Q_x, Q_y, G_{y,x}$ and G_x. Smaller size and special structure of $G_{y,x}$, G_x make them easy to update. However, since

$$\tilde{Y}\tilde{X}^T(\tilde{X}\tilde{X}^T)^{-1} = \tilde{Y}\tilde{X}^\dagger$$

this formulation increases the condition number $\kappa(\tilde{X}\tilde{X}^T) = \kappa(\tilde{X})^2$ making further computation potentially unstable. Additionally, since X and Y only differ up to one column, maintaining and updating two separate orthogonal matrices seems redundant.

Our approach is based on low-rank orthogonal decomposition of $S = Q\tilde{S}$ where $S = (s_1\ s_2\ \cdots\ s_{n+1})$. This way we only have to store and update one orthogonal basis Q. We propose an alternative formulation

$$\tilde{Y}\tilde{X}^{-1} = (\tilde{Y}\tilde{Q})T^{-1} =: G_{y,q}T^{-1}$$

using LQ or RQ decomposition of $\tilde{X} = T\tilde{Q}^T$ that avoids the squaring of the condition number while maintaining the same matrix sizes and simple updates. Our numerical analysis and numerical experiments demonstrate better numerical properties that result in better forecasting skill.

Domain decomposition method for eigenproblems

Author: Antti Hannukainen[1]

[1] *Aalto University*

Contact: antti.hannukainen@aalto.fi

I discuss solving large-scale symmetric and positive definite eigenproblems in distributed computing environments where communication between tasks is expensive, such as a cluster of networked workstations running the HTCondor batch system. As a model problem, I consider computing a few smallest eigenvalues of several eigenproblems related to FE-discretization of elliptic PDEs. The matrices are exported from COMSOL.

In particular, I develop a Ritz method using method subspace constructed from local spaces associated with non-overlapping partitions of the connectivity graph of the matrix. I discuss a technique for constructing these local subspaces from boundary-to-interior mapping related to the subdomain. I give an outline of error analysis and report numerical tests conducted on a cluster at Aalto University.

A-posteriori error estimates for randomized low-rank approximations

Authors: Lorenzo Lazzarino[1], Katherine J. Pearce[2], Nathaniel Pritchard[1]

[1] *University of Oxford*

[2] *University of Texas at Austin*

Contact: lazzarino@maths.ox.ac.uk

Randomized algorithms in numerical linear algebra have proven to be effective in ameliorating issues of scalability when working with large matrices, efficiently producing accurate low-rank approximations. A key remaining challenge, however, is to efficiently assess the approximation accuracy of randomized methods without additional expensive matrix accesses.

In this talk, we discuss a posteriori error estimation strategies for randomized low-rank approximations, with a focus on estimators that can be constructed from the same data used to compute the approximation, including leave-one-out type estimators. These can serve both as certification tools and as algorithmic building blocks, enabling adaptive approximations and informed trade-offs between accuracy and computational cost.

The Smith form of Sylvester and Bézout matrices for zero-dimensional ideals

Authors: Etna Lindy[1], Vanni Noferini[1]

[1] *Aalto University*

Contact: etna.lindy@aalto.fi

In the talk, I will go through our recent work regarding the Smith form of the Sylvester and Bézout resultant matrices. The partial multiplicities associated to the eigenvalue of a polynomial matrix tell us about the conditioning of computing the eigenvalue. Since the eigenvalues of the resultant matrices are the roots of the system, the partial multiplicities are connected to the stability of solving the polynomial system. We have also attempted to keep the assumptions as general as possible, and as a result, our coefficient field can be any field.

For bivariate polynomials $f, g \in \mathbb{K}[x, y]$, we make a connection between the partial multiplicities of the eigenvalues of the resultant matrices and the structure of the dual space of the zero-dimensional ideal $\langle f, g \rangle$. We show this by explicitly constructing the root polynomials for the matrix $S(y)$ using a Gauss basis of the dual space, introduced in [3M96].

In the talk, I aim to go through the general idea of the proof, as well as intuitively introduce the results through various examples. We first consider the case where f and g have coprime leading coefficients, as then $\det S(y) = \mathrm{Res}(f, g)$. Then, via means of a Möbius transformation that takes the "infinte roots" in the x-coordinate to finite roots, we prove that the Smith normal form stays intact through the transformation, allowing us to lift the earlier assumption.

[3M96] M. G. Marinari, H. M. Möller, T. Mora. On multiplicities in polynomial systems solving.

Trans. Amer. Math. Soc., 348(8):3283–3321, 1996.

A Superfast Direct Solver of Type-III Inverse Nonuniform Discrete Fourier Transform Problem

Authors: Jingyu Liu[1], Yingzhou Li[1]

[1] *Fudan University*

Contact: jyliu22@m.fudan.edu.cn

The nonuniform discrete Fourier transform (NUDFT) and its inverse are widely used in various fields of scientific computing. In this article, we propose a novel superfast direct inversion method for type-III NUDFT. The proposed method approximates the type-III NUDFT matrix as a product of a type-II NUDFT matrix and an HSS matrix, where the type-II NUDFT matrix is further decomposed into the product of an HSS matrix and an uniform discrete Fourier transform (DFT) matrix as in [Wilber, Epperly, and Barnett, SIAM Journal on Scientific Computing, 47(3):A1702-A1732, 2025]. This decomposition enables both the forward application and the backward inversion to be accomplished with quasi-linear complexity. The fast inversion can serve as a high-accuracy direct solver or as an efficient preconditioner. Additionally, we provide an error bound for the approximation under specific sample distributions. Numerical results are presented to verify the relevant theoretical properties and demonstrate the efficiency of the proposed methods.

The Polynomial Set Associated with a Fixed Number of Matrix-Matrix Multiplications

Authors: Elias Jarlebring[1], Gustaf Lorentzon[1]

[1] *KTH*

Contact: guslor@kth.se

We consider the problem of computing matrix polynomials $p(X)$, where X is a large dense matrix, using as few matrix-matrix multiplications as possible. Our approach to this problem involves studying the set $\Pi^*_{2^m}$, defined as the set of polynomials computable with at most m matrix-matrix multiplications, but with an arbitrary number of matrix additions and scaling operations. This set is characterized by a tabular parameterization. We derive equivalence transformations of the tabular parameterization, which can be used to identify and eliminate redundant parameters without changing the image space. This leads to two new scientific contributions: a) a precise determination of the dimension: $\dim \Pi^*_{2^m} = m^2$; b) a first minimal parameterization of the set. Using this parameterization, we develop new methods for constructing specific elements of $\Pi^*_{2^m}$. In addition, we use tools from computational algebraic geometry to determine the largest degree such that all polynomials of that degree belong to $\Pi^*_{2^m}$. We then use this knowledge to compute truncated

Taylor polynomials of maximal degree for a given number of matrix-matrix multiplications. This talk is based on the preprint [Jarlebring and Lorentzon, arXiv:2504.01500, 2025]

A Unified Framework for Linearizations of Matrix Polynomials in Classical Bases

Authors: Steve Mackey[1], Vasilije Perovic[2]

[1] *Western Michigan University*

[2] *University of Rhode Island*

Contact: steve.mackey@wmich.edu

Over the last two decades, a number of methods for constructing linearizations of matrix polynomials have been developed, including ansatz spaces, Fiedler pencils, and block minimal basis pencils. These methods have also been extended in various ways to apply to matrix polynomials expressed in non-monomial bases. However, these extensions have often been achieved one basis at a time, without systematically investigating the possibility of useful *relationships* between linearizations for matrix polynomials expressed in different bases. This talk describes a unified framework for establishing strong connections between generalized ansatz spaces associated with different bases, enabling the immediate and explicit construction of large spaces of strong linearizations for matrix polynomials expressed in many of the classical polynomial bases.

Fast and accurate solution of the periodic differential matrix Riccati equation

Author: Aleksandr Malyshev
Contact: alexander.malyshev@uib.no

We want to compute a T-periodic symmetric solution $X(t)$ of a T-periodic differential matrix Riccati equation

$$-\dot{X}(t) = X(t)A(t) + A^T(t)X(t) - X(t)B(t)R^{-1}(t)B^T(t)X(t) + Q(t)$$

such that all solutions of the feedback system $\dot{x} = [A(t) - B(t)R^{-1}(t)B^T(t)X(t)]x(t)$ are asymptotically stable, i.e. $\lim_{t\to\infty} x(t) = 0$. The standard solvability condition for the matrix Riccati equation requires that the T-periodic Hamiltonian matrix

$$H(t) = \begin{bmatrix} A(t) & -B(t)R^{-1}(t)B^T(t) \\ -Q(t) & -A^T(t) \end{bmatrix}$$

has no characteristic multipliers on the unit circle $|z| = 1$. The transition matrices $\Phi_k = \Phi_H(t_{k+1}, t_k)$ for the periodic system with the matrix $H(t)$ are calculated using the Runge-Kutta method for

$k = 0, 1, \ldots, P-1$ on an equidistant grid $t_k = kh$ with step $h = T/P$. The monodromy matrices for $H(t)$ sampled on the grid t_k are matrix products $\Psi_k = \Phi_{k+P-1} \ldots \Phi_{k+1} \Phi_k$, where the indices are taken modulo P. The sampled solution $X(t_k)$ is determined by the block eigenvalue problem

$$\Psi_k \begin{bmatrix} I \\ X(t_k) \end{bmatrix} = \begin{bmatrix} I \\ X(t_k) \end{bmatrix} \Lambda_k,$$

where the discrete-stable matrices Λ_k are the monodromy matrices of the aforementioned feedback system.

In paper [1], an orthogonal elimination method was proposed for accurate computation of Ψ_k. In paper [2], an $O(P \log P)$ algorithm was developed for computing all monodromy matrices Ψ_k. In this talk, I present a faster $O(P)$ method for computing the entire solution $X(t_k)$, $k = 0, 1, \ldots, P-1$. Alternative algorithms are discussed in paper [3].

References

[1] A.N. Malyshev, Siberian Math J., **30**(4):559-567, 1989.
[2] P. Benner et al, Parallel Algorithms for LQ Optimal Control of Discrete-Time Periodic Linear Systems, J. Parallel Distrib. Comp, **62**:306-325, 2002.
[3] S. Gusev et al, A numerical evaluation of solvers for the periodic Riccati differential equation, BIT Numer Math, **50**:301-329, 2010.

New Insights into the Equivalence of Thick and Implicit Restarting Lanczos

Authors: James Baglama[1], Kyle Monette[1], Vasilije Perovic[1]

[1] *University of Rhode Island*

Contact: kmonette@uri.edu

Given a symmetric tridiagonal matrix, it has been well–established, that in exact arithmetic, applying some of its eigenvalues as shifts via the QR strategy produces a particular structured matrix where the leading tridiagonal block contains the non–shifted eigenvalues and a trailing diagonal submatrix of the shifted eigenvalues. We will show that the leading tridiagonal block can be created, up to sign differences, by an orthogonal similarity reduction of a certain symmetric arrowhead matrix constructed directly from the non–shifted eigenvalues and associated eigenvectors. This also provides leverage to show a similar result with an upper bidiagonal matrix and a certain triangular arrowhead matrix. Both of the results set the foundation to show, in a unique way, that the well–known implicitly restarted Lanczos method is mathematically equivalent to thick-restarted Lanczos method with a similar extension to singular value problems using the Golub-Kahan Bidiagonalization process.

Polynomial Factorization via Matrix Representations

Authors: Azam Mozaffarikhah[1], Akuye-Shika Odametey[2], Isaac Agyei[2], Sergio R. López-Permouth[2]

[1] *Virginia Tech, Department of Mathematics*

[2] *Ohio University*

Contact: mozaffarikh@vt.edu

Polynomial factorization is classically studied within commutative polynomial rings, where irreducibility is an intrinsic algebraic property. In this talk, we present a linear-algebraic approach to factorization via entangled polynomial rings, in which polynomials are represented by structured matrices and analyzed using tools from matrix theory.

By embedding a polynomial into a family of $m \times m$ matrices (through m-nomials and their finite representations as Szab\'o rings), questions of reducibility are translated into questions about matrix factorizations. In this setting, reducibility is detected using the m-terminant, defined as the determinant of the associated matrix. Circulant matrices and their eigenvalues play a central role in determining when such factorizations occur.

We focus on the polynomials

$$f_p(x) = 1 + x + x^2 + \cdots + x^{p-1},$$

which are irreducible in $\mathbb{Q}[x]$ for prime p. Using determinant and eigenvalue computations, we show that while these polynomials remain irreducible for all $m < p$, their matrix representations become reducible at $m = p$. This leads to the notion of
emphvalence, the minimal matrix size at which reducibility occurs, and establishes that the valence of $f_p(x)$ is exactly p for all primes $p \geq 3$.

This perspective reframes polynomial factorization as a problem in linear algebra over rings, illustrating how matrix representations, determinants, and eigenvalues uncover hidden algebraic structure beyond the classical setting.

Convergence of flexible GMRES with and without randomized sketching

Authors: Lauri Nyman[1], John Pearson[2], Stefan Guettel[3]

[1] *University of Manchester*

[2] *University of Edinburgh*

[3] *The University of Manchester*

Contact: lauri.nyman@manchester.ac.uk

In this talk, we consider the theoretical convergence of flexible GMRES. While convergence of standard GMRES is well studied, there exist few results of similar nature for flexible GMRES. The aim of this talk is to discuss and fill in this gap. In addition, we report on experiences of using these ideas in the context of randomized sketching.

Randomized algorithms for operator trace estimation

Authors: Zvonimir Bujanovic[1], Luka Grubisic[2], Daniel Kressner[3], Hrvoje Olić[1]

[1] *Faculty of Science, University of Zagreb*

[2] *University of Zagreb, Faculty of Science, Department of Mathematics*

[3] *EPFL*

Contact: hrvoje.olic@math.hr

The implicit trace estimation is a problem of approximating the trace of a matrix A accessed only through matrix-vector multiplication $x \mapsto Ax$, with the goal of using as few multiplications as possible to obtain an accurate approximation. Girard-Hutchinson's estimator computes the ε-approximation using $\mathcal{O}(\varepsilon^{-2})$ products, while its' variance-reducing improvement, known as Hutch++, only requires $\mathcal{O}(\varepsilon^{-2})$ products. In this talk, we propose an extension of those algorithms to estimate the trace of a trace-class operator by using operator-function multiplications. Our estimator is based on the notion of a Gaussian distribution on an arbitrary Hilbert space, so precise high-probability error bounds can be derived analogously to the matrix case estimators. We compare the performance of our method with an existing algorithm of this type, the ContHutch++ algorithm by Zvonek, Horning & Townsend. The theoretical results we prove concern mostly positive definite operators, but we provide empirical results that show potential in applications to general trace-class operators.

On the Choice of Sign Defining Householder Transformations

Authors: Michael Overton[1], Pinze Yu[2]

[1] *New York University*

[2] *Yale University*

Contact: mo1@nyu.edu

It is well known that, when defining Householder transformations, the correct choice of sign in the standard formula is important to avoid cancellation and hence numerical instability. In this talk we point out that when the "wrong" choice of sign is used, the extent of the resulting instability depends in a somewhat subtle way on the data leading to cancellation.

Superfast 1-Norm Estimation

Authors: Soo Go[1], Victor Pan[1]

[1] *City University of New York*

Contact: soo.h.go@gmail.com

A matrix algorithm is said to be superfast (aka runs at sublinear cost) if it involves much fewer scalars and flops than an input matrix has entries. Such algorithms have been extensively studied and widely applied in modern computations for matrices with low displacement rank and more recently for low rank approximation of matrices, even though any superfast algorithm fails on worst case inputs for the latter problem. We extend this study to a new area by proposing three distinct superfast 1-norm estimators. In our extensive tests with synthetic and real word matrices all three have output 1-norm within relative error 1.5, while the outputs of two of them were consistently as accurate as LAPACK's. We point out some promising extensions of our surprisingly simple techniques. With further testing and refinement our algorithms should eventually win user's attention.

Superfast Low Rank Approximation

Authors: Victor Pan[1], Qi Luan[1], Soo Go[1]

[1] *City University of New York*

Contact: v_y_pan@yahoo.com

Low Rank Approximation (LRA) of a matrix are invaluable for Numerical Linear Algebra and Data Science. Some recent papers propose superfast algorithms that output LRAs with near-optimal accuracy for a large class of inputs but, as ANY superfast LRA algorithm, fail on a large class of inputs as well. To narrow the latter class we first superfast compute a crude initial LRA by applying one of these or another superfast algorithm (we specify a novel class of such algorithms). Then we recursively refine that LRA superfast by extending iterative refinement algorithms for the solution of the systems of linear and nonlinear equations and by applying the recent techniques of oversampling and compression of J. A. Tropp, A. Yurtsever, M. Udell, V. Cevher, Streaming Low-Rank Matrix Approximation with an Application to Scientific Simulation, SIAM J. on Scientific Computing, 41, pp. A2430–A2463, 2019; this enables us to control rank growth in the refinement. We analyze our algorithms and test them numerically.

Order automorphisms of effect algebras

Author: Peter Semrl[1]

[1] *Institute of Mathematics, Physics, and Mechanics, Ljubljana, Slovenia*

Contact: peter.semrl@fmf.uni-lj.si

The general form of order automorphisms of effect algebras has been known in the complex case. We present a much simpler proof based on projective geometry which works also in the real case.

Row and Column Equivalence Transversality Properties: Extensions of the Rank-Preserving Transversality Property

Authors: Alfred Hart, Frank Hall, Hanfei Xu, Hein Van Der Holst, Jiamin Pan, Jian Liu, Joonwon Seo[1], Li Wang, Marina Arav, Yiran Xu, Yue Zhao, Zheng Yang, Zhongshan Li, Ziuxan Li

[1] *Georgia State University*

Contact: jseo27@gsu.edu

This 20-minute contributed talk focuses on two properties introduced in recent work on matrix manifolds: the Row Equivalence Transversality Property (RETP) and the Column Equivalence Transversality Property (CETP). For an $m \times n$ real matrix A, RETP holds if the manifolds $\{GA : G \in \mathrm{GL}(m, \mathbb{R})\}$ and $Q(\mathrm{sgn}(A))$ intersect transversally at A, with tangent space $\{YA : Y \in \mathbb{R}^{m\times m}\}$. Similarly, CETP involves $\{AG : G \in \mathrm{GL}(n, \mathbb{R})\}$ with tangent space $\{AX : X \in \mathbb{R}^{n\times n}\}$. We establish that RETP and CETP each imply RPTP, but the converse does not hold (e.g., certain block matrices have RPTP without RETP/CETP). Key results include:

Theorem 4.1: A has RETP iff for each row with zeros, the corresponding columns are linearly independent.
Theorem 4.2: Dual for CETP with rows.
Dimension bounds: If RETP, then $m \cdot \mathrm{rank}(A) + \#(A) \geq mn$; similar for CETP.
Invariance under permutation and diagonal equivalence.
Theorem 4.4: RPTP with zero row/column implies full rank and RETP/CETP.
Corollaries for partitioning matrices to verify RPTP via RETP/CETP (e.g., Theorem 4.8, Corollary 4.9).
Examples: Matrices with RPTP but not RETP/CETP; conditions for both.

These properties aid in studying minimum ranks of sign and zero-nonzero patterns, with connections to transverse intersections in smooth manifolds. No prior RPTP knowledge assumed; proofs and examples provided.

A Classification of Seeds via the Minimum Number of Distinct Eigenvalues

Authors: Lucas Siviero Sibemberg[1], Carlos Hoppen[1], Lutz Emilio Allem[2]

[1] *UFRGS*

[2] *Universidade Federal do Rio Grande do Sul*

Contact: lucas.siviero@ufrgs.br

The underlying graph G of a symmetric matrix $M = (m_{ij}) \in \mathbb{R}^{n\times n}$ is the graph with vertex set $\{v_1, \ldots, v_n\}$ such that a pair $\{v_i, v_j\}$ with $i \neq j$ is an edge if and only if $m_{ij} \neq 0$. For a graph G, let $q(G)$ denote the minimum number of distinct eigenvalues among all symmetric matrices whose underlying graph is G. A symmetric matrix M is a realization of $q(G)$ if it has underlying graph G and exactly $q(G)$ distinct eigenvalues. For any tree T with diameter d, it is well known that $q(T) \geq d + 1$. We call T diminimal if equality holds, that is, if $q(T) = d + 1$.

Johnson and Saiago (2013) introduced a decomposition of the set of trees of any fixed diameter d based on a finite family of trees, called seeds, and on an operation known as combinatorial branch duplication (CBD, for short). A tree T' obtained from a tree T through a sequence of CBDs is called an unfolding of T. In particular, every tree with diameter d can be obtained as the unfolding of a unique seed of diameter d.

Relating the concept of seed with the minimum number of distinct eigenvalues problem, a seed S is called diminimal if every unfolding of S is diminimal, and defective otherwise. In this talk, we present ongoing results showing that most seeds are defective. In particular, the proportion of defective seeds tends to 1 as the diameter grows to infinity.

Graph Energy Change Under Edge Deletion

Authors: Wasin So[1], Huan-Lei Shi, Wen-Huan Wang

[1] *San Jose State University*

Contact: wasin.so@sjsu.edu

The energy $\mathcal{E}(G)$ of a graph G is the sum of the absolute values of all the eigenvalues of its adjacency matrix. Gutman (2001) proposed a problem of characterizing the graph G and the edge e of G such that $\mathcal{E}(G - e) \leq \mathcal{E}(G)$. Later, Day and So (2008) gave a sufficient condition : e is a cut-edge of G.Recently, Tang et al. (2023) gave another sufficient condition for $\mathcal{E}(G - e) \leq \mathcal{E}(G)$, where e is not a cut-edge.In this talk, we improve their result by giving a weaker sufficient condition and simplifying their proof. As a consequence, we prove that adding edges to a complete k-partite graph leads to higher energy than the energy of the original graph.

Tensor Train-Compressed FDTD Solvers for Electromagnetic Simulations

Authors: Erika Ye[1], Jonathan Tabares[2], Yang Liu[3]

[1] *Lawrence Berkeley National Laboratory*

[2] *Florida International University*

[3] *Lawrence Berkeley National Laboratory*

Contact: jtaba016@fiu.edu

This work presents a quantized tensor train (QTT)-accelerated finite-difference time-domain (FDTD) algorithm for solving Maxwell equations in 3D open domains.

Upon a QTT representation of both field variables and differential operators on uniform grids, the proposed algorithm can achieve up to logarithmic scaling in memory and per-step computational cost with respect to the number of spatial grid points. Our algorithm constitutes two major algorithmic contributions: (1) The integration of auxiliary differential equation perfectly matched layers (ADE–PML) directly in compressed form, enabling fully tensorized leapfrog time integration without decompression. (2) For complex excitations and media, QTT ranks can significantly grow due to both complex wave fields and numerical noise. To address this, we introduce spatial smoothing and filtering strategies that stabilize QTT ranks while preserving physical accuracy.

In summary, we demonstrate that classical explicit FDTD schemes can be fundamentally reinterpreted through QTT decompositions, providing new opportunities for scalable Maxwell solvers for large-scale and complex wave propagation simulations.

On some Graphs determined by their Signless Laplacian spectrum

Authors: Jitul Talukdar[1], Swarup Kumar Panda

[1] *IIT Kharagpur*

Contact: jitult00@gmail.com

In this paper, we investigate when the disjoint union of complete graphs $K_a \cup K_b$ is determined by its signless Laplacian spectrum (Q-DS). We first prove that $K_a \cup K_b$ is Q-DS among disconnected graphs. We then show that no connected signless Laplacian co-spectral mate of $K_a \cup K_b$ exists on at most ten vertices. Further, we establish that $K_t \cup K_a$ is Q-DS for $t = 1, 2$, except in the case $t = 1$ and $a = 3$. We also prove that $K_a \cup K_{a+1}$ is Q-DS. Finally, we show that the complete bipartite graph $K_{a,a+1}$ is determined by its signless Laplacian spectrum.

Newton's Method for Computing the CP Decomposition

Author: DuBose Tuller
Contact: tullwd25@wfu.edu

In this talk, I will address using Newton's Method to compute the CP tensor decomposition. The CP optimization problem is a nonlinear least squares problem with factor matrices as the variables. The most common methods for solving CP are Alternating Least Squares (ALS) and Gauss-Newton optimization combined with an iterative method for solving linear systems. I will show that one iteration of Newton's Method, when combined with an iterative linear solver, can achieve similar computational cost to ALS and Gauss-Newton for large, dense tensors. I will discuss the derivation for using Newton's method, how we achieve efficient computation, as well as experimental comparisons among these three methods.

Nearest correlation matrices with structure: a dynamical systems approach

Authors: Christian Lubich[1], Eugenio Turchet[2], Francesco Paolo Maiale[3], Nicola Guglielmi[3]

[1] *Universitaet Tuebingen*

[2] *GSSI*

[3] *Gran Sasso Science Institute*

Contact: eugenio.turchet@gssi.it

The nearest correlation matrix problem consists in finding the closest valid correlation matrix to a given symmetric matrix that may fail to be positive semi-definite. In other words, given a symmetric unit-diagonal matrix that is not a proper correlation matrix, one seeks the nearest positive semi-definite matrix with unit diagonal entries.

We address the problem of finding the nearest correlation matrix to a given symmetric unit-diagonal matrix under additional structural constraints such as sparsity, block, or band patterns. This task arises in applications where positive semi-definiteness must be restored without losing essential structure.

Our method combines a two-level iteration: a structured gradient flow computes feasible perturbations within the prescribed structure, while an outer Newton scheme adjusts their magnitude to meet accuracy requirements. To handle high-dimensional settings efficiently, we replace full eigenvalue decompositions with a Rayleigh quotient approximation, focusing only on the critical invariant subspace needed to restore positive semi-definiteness.

The resulting algorithm systematically incorporates structural constraints into the nearest correlation matrix problem. Numerical experiments highlight its robustness across diverse structured scenarios, with promising applications in finance, statistics, and network analysis.

References

1. N. J. Higham, *Computing the nearest correlation matrix — a problem from finance*, IMA J. Numer. Anal. **22** (2002), 329–343.

2. N. Guglielmi, C. Lubich, *The matrix nearness problem and eigenvalue optimization*, preprint, 2025.

3. N. J. Higham, N. Strabić, *Anderson acceleration of the alternating projections method for computing the nearest correlation matrix*, Numer. Algorithms **72** (2016), no. 4, 1021–1042.

Fast computation of eigenvalues of periodic CMV matrices

Author: David S. Watkins[1]

[1] *Washington State University*

Contact: davidwatkinsmatrix@gmail.com

Periodic CMV matrices are unitary matrices that can be specified by $O(n)$ data. Their eigenvalues can be computed by standard methods, storing them as conventional matrices (using $O(n^2)$ data) in $O(n^3)$ time. Since periodic CMV matrices can be specified by $O(n)$ data, one would hope to find a method that computes the eigenvalues in $O(n^2)$ time instead of $O(n^3)$. This is indeed possible, and we will show how to do it.

On the properties of solution sets of absolute value equations

Authors: Dipti Dubey[1], Shweta Yadav[1]

[1] *Shiv Nadar University*

Contact: sy213@snu.edu.in

We investigate the structural properties of the solution set of absolute value equations (AVE) of the form $Ax - |x| = b$. Extending the seminal work of Hladík (SIAM J. Matrix Anal. Appl., 2023), we address his open questions originally posed for $Ax + |x| = b$ and establish analogous results for the alternative form considered here. Using the equivalence between AVE and the linear complementarity problem (LCP), we derive new results on the convexity and solvability of the AVE under various matrix classes, including positive semidefinite, Metzler, and Z-matrices. Further, we present sufficient conditions guaranteeing unique and nonnegative solutions for $b \geq 0$ or $b \leq 0$, and provide counterexamples showing that several properties valid for $Ax + |x| = b$ fail to extend to $Ax - |x| = b$. These findings contribute to the theoretical understanding of AVEs and enrich the connection between matrix theory and complementarity problems.

Householder Sets and Relative Pseudospectra

Authors: Paul Zachlin[1], Michiel Hochstenbach[2]

[1] *Lakeland Community College*

[2] *TU Eindhoven*

Contact: pzachlin@lakelandcc.edu

We study eigenvalue inclusion regions first defined by Householder and explore applications to large, sparse matrices. We also study pseudospectra and relative pseudospectra, which are similarly defined eigenvalue inclusion regions. Then we explore connections between these sets.

List of Speakers & Contributions

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X

Y

Z