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

Monday, May 18, 2026 - Friday, May 22, 2026

Virginia Tech

Scientific Program

Plenary Talks

Minisymposia Sessions

Numerical Linear Algebra in Machine Learning

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.

New Directions and Challenges in Linear Algebra

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.

Advanced Acceleration and Convergence Techniques for Solving Linear and Nonlinear Systems

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.

Matrix Geometries

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.

Matrix Nearness Problems

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.

Matrix Inequalities, Matrix Equations, and Their Applications

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.

Spectral Graph Theory

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.

The Inverse Eigenvalue Problem of a Graph and Zero Forcing

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.

Numerical Ranges and Numerical Radii

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.

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

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.

Numerical Linear Algebra Tools for Model Order Reduction

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.

Theoretical Advances in Operator Learning

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.

Computational Advances in Discrete Inverse Problems

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.

Code-based Cryptography

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.

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

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.

Hierarchical Low-Rank Approximations: Algorithms and Applications

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.

Where Algebraic Coding Theory and Graph Theory Meet

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.

Topics in Randomized Numerical Linear Algebra

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.

Convex Structures in Quantum Information and Gravity

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.

Eigenvalues of Nonnegative and Stochastic Matrices

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.

Krylov Iterative Methods for Linear Equations

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.

Algebraic Invariants of Graphs

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.

Low-Complexity Data-driven or Classical Algorithms and Applications

Organizer: 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 non-linear 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.

Approximate Computing in Numerical Linear Algebra

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.

Rational Approximation and Interpolation: Practical Applications, Challenges and Solutions

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.

Advances and Challenges in Eigensolvers

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.

Linear Algebra Education

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.

Symplectic Linear Algebra and Applications

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.

Low-rank Matrix and Tensor Decompositions: Theory, Algorithms and Applications

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.

Spectral Interlacing, Graph Learning, and Quantum Perspectives on Signed Graphs

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.

Sparse Tensor Computations: Algorithms and Applications

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.

Advances in Randomized Algorithms and Kernel Methods for Rank-Structured Matrices

Organizers: Rajarshi Bhattacharjee, Tong Ding, Vladimir Druskin, Marc Aurele Gilles, Mikhail Lepilov, Siting Liu, Ibrohim Nosirov, 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.

Quantum Numerical Linear Algebra

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.

Model- and Data-driven Reduced-order Models and Their Applications in Inverse Problems

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.

Recent Advances in Tensor Decompositions for Model and Data Reduction

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.

Polynomials, Krylov Methods and Applications

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.

New Advancements in Tensor Decomposition and Computation

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.

Combinatorial Matrix Theory

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.

Advances in Application-Driven Family of Matrix Computations: Factorization, Inverse, and Linear Solves

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.

Contributed Talks