YMMOR2026

program & abstracts

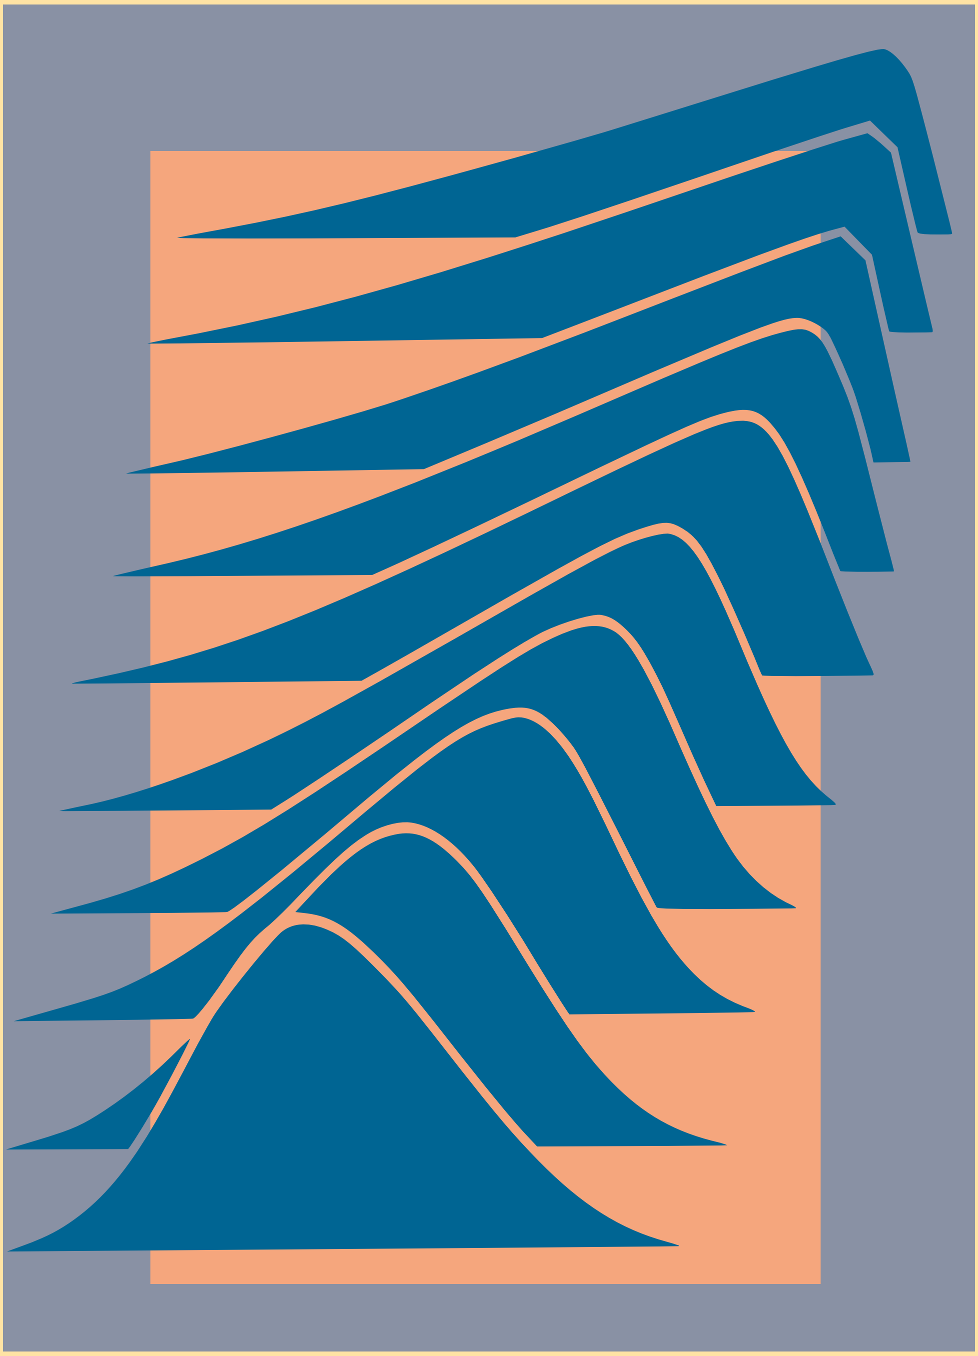

Thank you to our sponsors!

COLLEGE OF SCIENCE
MATHEMATICS
VIRGINIA TECH

COLLEGE OF SCIENCE
VIRGINIA TECH

COLLEGE OF SCIENCE
ACADEMY OF DATA SCIENCE
VIRGINIA TECH

HYATT

Conference Program

Time Slot	Tuesday	Wednesday	Thursday	Friday
Morning Coffee (9:00AM-9:30AM)				
9:30AM-11:30AM	Aryeah Keating	Ziqi Huang	Rudi Smith	Cankat Tilki
	Gregory Janes	Jakob Scheffels	Rodrigo Figueroa	Benedikt Klein
	Tomoki Koike	Weiting Yi	Leonie Van Pottelberghe	Robin Herkert
	Ela Đimoti	Abhinab Bhattacharjee		Jenifer De Jager
Lunch Break (11:30AM-1:00PM)				
1:00PM-2:30PM	Mike Ackermann	Amy de Castro	Emmanuel Ameh	Closing remarks
	Albani Olivieri	Andrew Beusse	Edward Huynh	
	Ian Moore	Annamaria Palminiero	Sven Ullmann	
2:30PM-3:30PM	Breakout Session 1	Breakout Session 2	Breakout Session 3	
Coffee Break (3:30PM-4:00PM)				
4:00PM-5:00PM	Sam Bender	Jordan Jackson	Dan Folescu	
	Till Peters	Hayden Ringer	Francesco A.B. Silva	

Length of talks Each speaker is alotted **30** minutes for their presentation, **including questions**.

Group excursion (Monday, 10:00AM-2:00PM) We will travel to and hike along the **Cascades Fall Trail**. This a moderate, 3.7 mile (5.9 km) long hike that will take you through some of the most gorgeous scenery of the Appalachians to a 69-foot-tall waterfall.

Welcome picnic (Monday, 2:00PM-5:00PM) Following the group excursion, a welcome picnic will be held at the Blacksburg Municipal Park, located at 615 Patrick Henry Drive, in Shelter 9.

Conference dinner (Wednesday, 6:00PM-8:30PM) The conference dinner will be hosted at **The Maroon Door**, located in the heart of downtown Blacksburg at 418 North Main Street.

Book of Abstracts

From data to structured models via second-order AAA algorithms

M. S. Ackermann[1], I. V. Gosea[2], S. Gugercin[3], and S. W. R. Werner[4]

[1]*Department of Mathematics, Virginia Tech*
[3]*Department of Mathematics and Division of Computational Modeling and Data Analytics, Virginia Tech*
[4]*Department of Mathematics, Division of Computational Modeling and Data Analytics, and National Security Institute, Virginia Tech*
[2]*Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany*

Constructing dynamical system models directly from real-world measurements has become an essential tool in modern science and engineering. When the internal physics of the considered process are known, the derivation of dynamical system models from first principles frequently yields internal differential structures. For example the system may exhibit an explicit dependence on the second time-derivative

$$M\ddot{x}(t) + D\dot{x}(t) + Kx(t) = bu(t), \quad y(t) = c^T x(t),$$

with system matrices $M, D, K \in \mathbb{R}^{n\times n}$, $b, c \in \mathbb{R}^n$, the inputs $u(t) \in \mathbb{R}$, outputs $y(t) \in \mathbb{R}$, and internal state variable $x \in \mathbb{R}^n$. The matrix M is the mass matrix, D is the damping matrix, and K is the stiffness matrix. As suggested by their names, these matrices are important for interpreting properties of the dynamical system model.

We are interested in producing second-order structured dynamical systems from frequency-domain data, which is the common format for data in mechanical and acoustic modeling applications. Data-driven methods are able to construct highly accurate but unstructured surrogate models when explicit state-space models are not available, but input-output data is abundant. The adaptive Anderson Antoulas (AAA) algorithm [4] is one such data-driven modeling algorithm, which produces unstructured models from frequency domain data by blending interpolation and least-squares fitting. Based on our previous work with structured barycentric forms [3], we derive second-order structured adaptive Anderson Antoulas (AAA) algorithms, each of which learn second-order structured models from data. Enforcing the second order structure introduces additional nonlinearities into the optimization problem, which we carefully treat to ensure a robust algorithm. We present the theoretical foundations of these algorithms and examplify their effectiveness for structured data-driven modeling through numerical examples.

A preprint for this work is available [1], as well as code to reproduce all experiments [2].

References

[1] M. S. Ackermann, I. V. Gosea, S. Gugercin, and S. W. R. Werner. Second-order aaa algorithms for structured data-driven modeling, 2025.

[2] M. S. Ackermann and S. W. R. Werner. Code, data and results for numerical experiments in "Second-order AAA algorithms for structured data-driven modeling" (version 1.0), June 2025.

[3] I. V. Gosea, S. Gugercin, and S. W. R. Werner. Structured barycentric forms for interpolation-based data-driven reduced modeling of second-order systems. *Adv. Comput. Math.*, 50(2):26, 2024.

[4] Y. Nakatsukasa, O. Sète, and L. N. Trefethen. The AAA algorithm for rational approximation. *SIAM J. Sci. Comput.*, 40(3):A1494–A1522, 2018.

Trajectory-Optimization-Based Model Reduction via Balancing States and Costates for Optimal Control

A. Emmanuel Sunday[1] **and O. Samuel**[1]

[1]*Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY, USA*

Robust and 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 drag reduction and recovering from aerodynamic stall. With initial conditions far away from the desired steady state, accounting for nonlinear effects becomes crucial. However, this makes solving the resulting Hamilton-Jacobi-Bellman equation (HJB), which defines the value function over the continuous state space, crucial for optimal feedback control, intractable as the computational complexity scales exponentially with the state-space dimension under discretization. ROMs can help in a variety of ways, but existing methods often fail to capture relevant dynamics for the control problem [1].

To overcome these challenges, we first approximate the infinite-horizon optimal control problem with a large finite horizon to obtain a time-invariant version of the value function. We furthermore employ an indirect method of gradient-based trajectory optimization (which is feasible in high dimensions) to obtain state and costate data along locally optimal trajectories and estimate state and gradient covariance matrices. This method builds on the Pontryagin minimum principle and other related work that establishes the costate (adjoint variables) provided by an adjoint ODE solved backward in time as generalized gradients of the optimal value function [2]. With initial conditions sampled uniformly from an uncontrolled attractor, such as a limit cycle in fluid flows that undergo Hopf bifurcations or other user-specified distribution, these matrices estimated by Monte Carlo sampling are balanced and used to identify an active subspace for the high-dimensional value function. The balanced active-subspace identifies directions in the state space along which the value function is most sensitive, and states have large variance [3].

An oblique projection obtained by balancing is used to build polynomial surrogate models for both the value function and the optimal feedback control law, which is validated on the full order model. We also provide error estimates for the approximate value function in the active subspace. This methodology is tested on benchmark fluid dynamical systems, including the complex Ginzburg-Landau and Kuramoto-Sivashinsky equations, and is compared against benchmark ROM techniques.

References

[1] A. Alla, A. Schmidt, and B. Haasdonk. Model order reduction approaches for infinite horizon optimal control problems via the hjb equation. In P. Benner, M. Ohlberger, A. Patera, and G. Rozza, editors, *Model Reduction of Parametrized Systems*, pages 333–347. Springer, Cham, 2017.

[2] F. H. Clarke and R. B. Vinter. The relationship between the maximum principle and dynamic programming. *SIAM Journal on Control and Optimization*, 25(6):1291–1311, 1987.

[3] S. E. Otto, A. Padovan, and C. W. Rowley. Model reduction for nonlinear systems by balanced truncation of state and gradient covariance. *SIAM Journal on Scientific Computing*, 45(5):A2325–A2355, 2023.

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

S. Bender[1] **and C. Beattie**[1]

[1]*Department of Mathematics, Virginia Tech, Blacksburg, 24061, VA, United States*

We consider single input, single output systems with periodic parameters:

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{b}(t)u(t), \quad y(t) = \mathbf{c}(t)^*\mathbf{x}(t), \tag{1}$$

where $\mathbf{A}(t) \in \mathbb{R}^{n \times n}$, $\mathbf{b}(t), \mathbf{c}(t) \in \mathbb{R}^n$ all have period T. Such systems arise as the result of modeling phenomena related to fluid dynamics, structural mechanics, and electronic circuits. Specifically, linearization around known periodic orbits of a nonlinear model produces a periodic system of partial differential equations, then semi-discretization in space yields large scale linear time-periodic (LTP) dynamical systems. The need to simulate responses to a variety of inputs motivates the development of effective model reduction tools for these systems.

While the research on model reduction for LTP systems is limited, there is a sizeable amount of literature devoted to control, spectral analysis, and harmonic response of LTP systems [1, 2, 3]. Arising from this literature is the Harmonic Transfer Function and the $\mathcal{H}_2$ norm for LTP systems. These conceptions lead to necessary conditions for an $\mathcal{H}_2$ optimal reduced order model. In the linear time-invariant case, the Iterative Rational Krylov Algorithm (IRKA) [4] is a standard approach for the $\mathcal{H}_2$ model reduction problem. This work presents advancements towards adapting the IRKA to the LTP setting.

References

[1] H. Sandberg, E. Mollerstedt, and Bernhardsson. Frequency-domain analysis of linear time-periodic systems. *IEEE Transactions on Automatic Control*, 50(12):1971–1983, December 2005.

[2] N.M. Wereley and S.R. Hall. Frequency response of linear time periodic systems. In *29th IEEE Conference on Decision and Control*, pages 3650–3655 vol.6, 1990.

[3] Jun Zhou and Tomomichi Hagiwara. H_2 and H_∞ norm computations of linear continuous-time periodic systems via the skew analysis of frequency response operators. *Automatica*, 38(8):1381–1387, August 2002.

[4] S. Gugercin, A. Antoulas, and Christopher Beattie. H_2 model reduction for large-scale linear dynamical systems. *SIAM Journal on Matrix Analysis and Applications*, 30:609–638, 06 2008.

Modal Decomposition of Boundary Layer Flow Over Diverse Walls

A. Beusse[1], A. D. G. Hales[1], R. Meadows[1], and S. Damani[1]

[1] *Virginia Polytechnic Institute and State University*

Wall-bounded flows are of significant interest to the aerospace engineering community and are popular flows studied by fluid dynamicists alike. For wall-parallel turbulent flow, the presence of large-scale **superstructures** have long intrigued the turbulence research community due to their dominant role in velocity and pressure fluctuations. These structures, characterized by their long streamwise meandering nature and prominence at low frequencies, have been extensively observed in turbulent boundary layers over both smooth and rough surfaces [2, 6] . Typically residing in the logarithmic region, these superstructures transport momentum and contribute significantly to turbulent energy transfer [4, 3]. Naturally, the presence of physically significant coherent structures within the data appeals to data-driven methods like model reduction to extract, quantify and understand the role they play within the flow and beyond.

In this talk, we outline past and present research performed by the Devenport lab to apply numerous **DMD algorithms** to reliably extract modes responsible for these features, as performed in [1, 5]. These algorithms are chosen to provide verifiable modes that can generate sensible synthetic flows to characterise.

We shed light on how these methods can help distinguish and characterise differences in turbulent flow over homogeneous smooth and rough walls, as well as heterogenous smooth-to-rough walls.

References

[1] H. Butt, S. Damani, S. Srivastava, S. S. Chaware, M. Szoke, W. J. Devenport, T. Lowe, A. Hales, M. Colbrook, and L. J. Ayton. Pressure gradient effects on boundary layer superstructures. In *AIAA AVIATION 2023 Forum*, page 3337. American Institute of Aeronautics and Astronautics, 2023.

[2] H. Butt, B. Sharma, S. Damani, E. Totten, T. K. Lowe, and W. J. Devenport. Roughness impacts on boundary layer super structures. In *13th International Symposium on Turbulence and Shear Flow Phenomena (TSFP13)*, 2024.

[3] B. Ganapathisubramani, N. Hutchins, W. Hambleton, E. Longmire, and I. Marusic. Investigation of large-scale coherence in a turbulent boundary layer using two-point correlations. *Journal of Fluid Mechanics*, 524:57–80, 2005.

[4] N. Hutchins and I. Marusic. Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *Journal of Fluid Mechanics*, 579:1–28, 2007.

[5] B. Sharma, A. Hales, H. Butt, S. Damani, E. D. Totten, W. J. Devenport, T. Lowe, and M. Colbrook. Characterizing superstructures in rough-wall turbulence. In *AIAA AVIATION FORUM AND ASCEND 2025*.

[6] Y. Wu and K. T. Christensen. Spatial structure of a turbulent boundary layer with irregular surface roughness. *Journal of Fluid Mechanics*, 655:380–418, 2010.

Multi-Fidelity Particle Flow Filter

A. Bhattacharjee[1], **A. Popov**[2], **A. Subramanya**[3], **and A. Sandu**[1]

[1] *Computational Science Lab, Department of Computer Science, Virginia Tech, USA*
[2] *Department of Information & Computer Sciences, the University of Hawai'i at Manoa, Honolulu, Hawai'i, USA*
[3] *Argonne National Laboratory, Lemont, Illinois, USA*

This work proposes a new Multi Fidelity Particle Flow Filter (MFPFF) based on a linear control variate framework and multilevel Monte Carlo methods. Particle flow filters "drifts" the prior particles towards the most likely posterior particles using an advection law. A diffusion is often added to keep the particles well spread out. Particle flow filters, in general, avoid the particle degeneracy, as is often a problem in standard particle filters by bypassing pointwise multiplication of two functions. We explore two-fidelity data assimilation with both stochastic (Gromov) and deterministic (exact) flow in the well-known Korteweg-De Vries (KdV) and Quasi-Geostrophic (QG) test problems. In this setup, we augment a few full-order expensive model runs with multiple reduced-order inexpensive surrogates to have a better estimate compared to a simulation with only full-order model. We use an intrusive Galerkin ROM as our cheap surrogates. This algorithm can be extended to multi-fidelity setting. The numerical experiments support our claim.

Reduced order modeling for fluid interaction systems with elastic or poroelastic structures

Amy de Castro[1] **and Hyesuk Lee**[2]

[1] *University of Utah*
[2] *Clemson University*

In this talk, we discuss the implementation of projection-based reduced order models (ROMs) into a strongly coupled partitioned method for the solution of fluid-structure interaction problems with linear elastic [2, 3] or poroelastic [1] structures. The partitioned scheme centers around a Schur complement equation whose solution is the fluid pressure and Lagrange multipliers representing interfacial quantities. At each time step, solving the Schur complement equation effectively decouples the physical subdomains and allows for their independent update without requiring iterations between the subdomains.

To reduce computational costs and to provide a more robust framework, we focus on the introduction of ROMs into this scheme. Utilizing the supremizer enrichment technique, we investigate the performance of this method with respect to the use of supremizers and with respect to the basis sizes of the reduced order variables. In both the reproductive and predictive regimes, results indicate that the ROM-ROM coupled formulation yields results that agree well with the full order solution in a shorter computational time and with a large reduction in the size of the algebraic system.

References

[1] A. de Castro and H. Lee. Well–posedness of a novel Lagrange multiplier formulation for fluid–poroelastic interaction. In submission. Preprint available at https://arxiv.org/abs/2512.08142, 2025.

[2] A. de Castro, H. Lee, and M. Wiecek. Reduced order modeling for a Schur complement method for fluid–structure interaction. *Journal of Computational Physics*, 515:113282, 2024.

[3] A. de Castro, H. Lee, and M. Wiecek. A Lagrange multiplier method for fluid–structure interaction: Well–posedness and domain decomposition. *Computers & Mathematics with Applications*, 181:193–215, 2025.

Lebesgue Boundedness of the Koopman Operator

J. De Jager[1] **and S. Gugercin**[2]

[1,2] *Department of Mathematics, Virginia Tech*

A number of data-driven model order reduction algorithms, including Extended Dynamic Mode Decomposition [3], make use of the Koopman operator, that is the (left) composition operator of the function defining the dynamics, acting on $\mathcal{L}^2$ for their analysis.
This operator was first introduced by Koopman in [1] and has seen widespread use in ergodic theory, see, e.g., [2], where it is referred to as the associated operator. In ergodic theory, the systems are measure preserving, which results in the Koopman operator being unitary ([2], Lemma 2.18).
In the context of EDMD, however, one may not be working in a measure preserving context, which complicates the boundedness and related properties of the Koopman operator. In this talk, we will discuss a framework for its boundedness and derive a sharp bound on the norm of the Koopman operator over the $\mathcal{L}^2$ space associated with the Lebesgue measure.

References

[1] B. O. Koopman. Hamiltonian systems and transformations in hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, 17(5):315–318, 1931.

[2] T. W. Manfred Einsiedler. *Ergodic Theory with a view towards Number Theory.* Springer, 2011.

[3] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data–driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, jun 2015.

Streaming Dynamic Mode Decomposition

E. Đimoti[1], joint work with Z. Drmač[1]

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

Dynamic Mode Decomposition (DMD) is a data-driven tool for capturing complex nonlinear dynamics. It can be used to identify, analyze and forecast continuous $\dot{x} = T(x)$ or discrete dynamical systems $x_{i+1} = T(x_i)$, governed by a mapping $T : \mathbb{R}^m \longrightarrow \mathbb{R}^m$ that is either unknown or too complex for direct analysis. It is grounded in the Koopman operator framework which allows for the representation of a nonlinear dynamical system via a linear, but infinite-dimensional operator $\mathcal{K}$. For simplicity we denote observed snapshots by $(x_i, y_i) := (x_i, T(x_i))$ and form data matrices $X := \begin{pmatrix} x_1 & x_2 & \cdots & x_n \end{pmatrix}$, $Y = \begin{pmatrix} y_1 & y_2 & \cdots & y_n \end{pmatrix}$. It can be shown that the finite-dimensional approximation of the Koopman operator is an $m \times m$ matrix $\mathbb{A}$ such that

$$\mathbb{A}X \approx Y$$

in the least-squares sense, i.e. $\mathbb{A}$ is the solution to $\min_A ||AX - Y||_F$. The two main tasks of DMD then are to

1) Identify approximate eigenpairs (λ_i, z_i) such that

$$\mathbb{A}z_i \approx \lambda_i z_i$$

2) Derive a spatio-temporal representation of snapshots y_i using selected eigenpairs $(\lambda_{\omega_1}, z_{\omega_1})$,..., $(\lambda_{\omega_l}, z_{\omega_l})$

$$y_i \approx \sum_{j=1}^{l} z_{\omega_j} \alpha_j \lambda_{\omega_j}^i \tag{1}$$

Key features of the DMD are that it is entirely data driven, oblivious to the nature of the data, and it offers a spatio-temporal decomposition representation of snapshots. The spatio-temporal representation of snapshots at the same time reveals and reduces the original model with unknown dynamics to the model obtained by (1) that uses only $l < \min(m, n)$ eigenpairs and often $m \gg n$.

In an online or streaming application, data snapshots are received in discrete time steps, possibly in batches, so that the DMD has to be recomputed over a widening data window. At some point, the oldest data may gradually be forgotten. Further, to capture sudden changes in the dynamics and keep the forecasting skill in case of a black swan event, the oldest data are immediately discarded, which means that the data window may suddenly shrink and then keep widening again. All this requires fast updating of the decomposition, that matches the dynamics of the data windows widths, and restoring the forecasting capabilities. We revisit the work of Hemati et al. [2] and Zhang et al. [3] on the online DMD and propose a more robust approach based on (low-rank) orthogonal decomposition. Our numerical analysis and numerical experiments [1] demonstrate better numerical properties that result in better forecasting skill. From the software development point of view, we define a modular framework for efficient software development, porting and maintenance on high performance hardware and software platforms.

References

[1] Z. Drmač and E. Đimoti. Streaming DMD and dynamic forecasting. Technical report, Department of Mathematics, Faculty of Science, University of Zagreb, Croatia, 2026.

[2] M. S. Hemati, M. O. Williams, and C. W. Rowley. Dynamic mode decomposition for large and streaming datasets. *Physics of Fluids*, 26(11), 2014.

[3] H. Zhang, C. W. Rowley, E. A. Deem, and L. N. Cattafesta. Online dynamic mode decomposition for time-varying systems. *SIAM Journal on Applied Dynamical Systems*, 18(3):1586–1609, 2019.

Block Structure Tensor Sketching for High Dimensional Problems

Paul Cazeaux[1], **Mi-Song Dupuy**[2], **and Rodrigo Figueroa**[1]

[1] *Virginia Tech*
[2] *Sorbonne Universite*

Large-scale scientific computing has render new bottlenecks that make computation infeasible. Randomized numerical linear algebra addresses this challenge through sketching, i.e., random dimensional reduction preserving essential geometric structure, while tensor-network formats such as the tensor-train (TT) decomposition provide compact representations of exponentially large multiway arrays, with impactful applications in quantum physics, parametric PDEs, and modern machine learning. However, existing tensor sketching methods either exhibit exponential scaling in tensor dimension or lack rigorous theoretical guarantees.

We propose a novel block-structured sketching framework formed by a concatenated sequence of small sketches aligned with TT cores, thus avoiding large Kronecker products or dense operators, while achieving provable embedding guarantees, natural parallelization, streamability, and scalability, and enabling efficient algorithms for randomized compression, Krylov eigensolvers, and streaming approximation in tensor formats.

Contour Integral Methods for the Masses

Dan Folescu[1], **Mark Embree**[1], **and Serkan Gugercin**[1]

[1] *Virginia Tech*

Linear and nonlinear eigenvalue problems alike can be solved using contour integral methods. These methods can be framed in the language of data-driven system identification [1]; structured quadrature data and problem-dependent hyper-parameters are compiled into data matrices, whereby local spectral information can be extracted through system realization. As the construction of these data matrices typically results in ill-conditioned Hankel/Loewner matrix pencils, various approaches can be taken to improve the stability of the system realization step.

`https://github.com/dan123222123/CIMTOOL`

In this talk, we will compare and contrast some of these approaches utilizing a new MATLAB tool, which facilitates the construction of the aforementioned data matrices from both exact and inexact data obtained via numerical integration. Self-consistent and visually reactive code structures allow for real-time exploration and refinement of method parameters, either programmatically or through the user-friendly graphical interface. Various numerical examples will showcase how to get the most out of contour integral methods and this new tool. This research was funded by US National Science Foundation Grant DMS-2411141.

References

[1] M. C. Brennan, M. Embree, and S. Gugercin. Contour Integral Methods for Nonlinear Eigenvalue Problems: A Systems Theoretic Approach. *SIAM Review*, 65(2):439–470, May 2023.

Randomized Symplectic Model Order Reduction

Robin Herkert[1], **Patrick Buchfink**[1], **Bernard Haasdonk**[1], **Johannes Rettberg**[2], **and Jörg Fehr**[2]

[1]*Institute of Applied Analysis and Numerical Simulation, University of Stuttgart, Stuttgart, Germany*
[2]*Institute of Engineering and Computational Mechanics, University of Stuttgart, Stuttgart, Germany*

We consider large-scale Hamiltonian systems, which in canonical coordinates $x = (q, p) \in \mathbb{R}^{2n}$ admit the formulation

$$\dot{x} = J\nabla H(x), \qquad J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}, \tag{1}$$

for a Hamiltonian $H \in C^1(\mathbb{R}^{2n}, \mathbb{R})$. For multi-query and real-time applications, projection-based model order reduction (MOR) is indispensable, but generic reductions may not preserve the symplectic geometry leading to long-time errors and energy gain/loss. Structure-preserving MOR therefore seeks reduced spaces that are symplectic so that the reduced dynamics remains Hamiltonian which ensures energy conservation.

Symplectic bases are commonly extracted from snapshot data via the complex singular value decomposition (cSVD) or related SVD-like decompositions. When the full dimension and/or the snapshot count is large, these factorizations can dominate the offline cost. This contribution summarizes randomized symplectic basis generation methods that combine randomized sketching with structure-preserving postprocessing, in particular the randomized complex SVD (rcSVD) and a randomized SVD-like decomposition, yielding substantial speedups in basis construction [1].

Beyond computational efficiency, we highlight recent theoretical advances that provide *a priori* error bounds for rcSVD depending on key hyperparameters (e.g., sketch size/oversampling and power-iteration depth). These results quantify the trade-off between runtime and approximation quality and are supported by numerical experiments demonstrating that randomized symplectic MOR can achieve accuracy close to cSVD at reduced cost [2].

Keywords: Hamiltonian systems; symplectic model order reduction; randomized algorithms; error bounds.

References

[1] R. Herkert, P. Buchfink, B. Haasdonk, J. Rettberg, and J. Fehr. Randomized symplectic model order reduction for hamiltonian systems. In *Large-Scale Scientific Computations (LSSC 2023)*, pages 99–107. Springer, Cham, 2024.

[2] R. Herkert, P. Buchfink, B. Haasdonk, J. Rettberg, and J. Fehr. Error analysis of randomized symplectic model order reduction for hamiltonian systems. *Linear Algebra and its Applications*, 729:67–99, 2026.

Error Analysis of Bayesian Inverse Problems with Generative Priors

B. Hosseini[1] **and Z. Huang**[2]

[1,2]*University of Washington, Seattle, WA 98195, USA*

Abstract Data-driven methods for the solution of inverse problems have become widely popular in recent years thanks to the rise of machine learning techniques. A popular approach concerns the training of a generative model on additional data to learn a bespoke prior for the problem at hand. In this article we present an analysis for such problems by presenting quantitative error bounds for minimum Wasserstein-2 generative models for the prior. We show that under some assumptions, the error in the posterior due to the generative prior, will inherit the same rate as the prior with respect to the Wasserstein-1 distance. We further present numerical experiments that verify that aspects of our error analysis manifests in some benchmarks followed by an elliptic PDE inverse problem where a generative prior is used to model a non-stationary field.

Summary of Contributions Below we summarize the main theoretical and numerical contributions of this work.

- *Main Theory:* We extended the theory on [1] by considering any map $\widehat{T} \in \arg\min_{T \in \widehat{\mathcal{T}}} \mathcal{W}_2(T\#\eta, \mu^N)$, where $\widehat{\mathcal{T}}$ denotes an approximation class such as neural nets or polynomials, and μ^N denotes the empirical measure associated with N i.i.d. samples from μ. Under appropriate assumptions, we prove that for any $\varepsilon > 0$ and with probability $1 - CN^{-1/d}/\varepsilon$ (with some constant $C > 0$) it holds that
$$\mathcal{W}_1(\widehat{\nu}, \nu) \lesssim \inf_{T \in \widehat{\mathcal{T}}} \|T - T^\dagger\|_{L^2_\eta} + \varepsilon + \delta$$
where ν is the true posterior and $\widehat{\nu}$ an approximate posterior measure that arises due to an approximate prior $\widehat{\mu}$. $\lesssim$ contains independent constants, $T^\dagger$ is any map that satisfies $T^\dagger\#\eta = \mu$ and $\delta > 0$ is an additional bias term that encodes the tail properties of the prior μ.

- *Numerical Results:* We present a number of numerical experiments aimed at verifying our theoretical bounds, Figure 1 shows representative examples of generative priors and sampled posteriors in a two-dimensional setting.

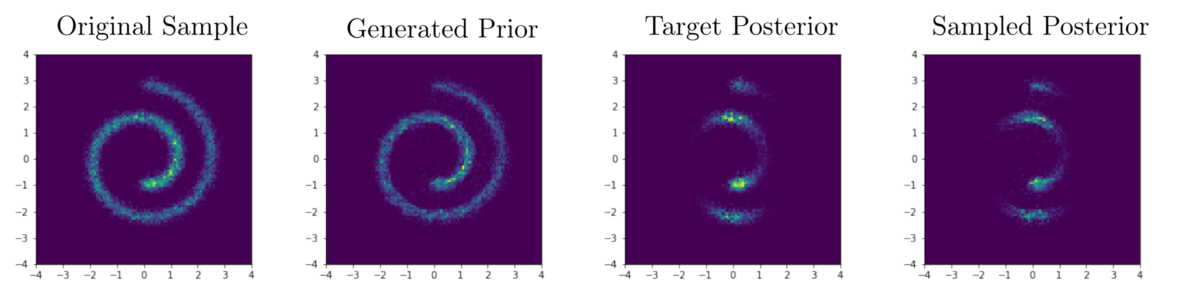

Figure 1: A 2D example: From left to right are the samples of original distribution, generated prior distribution, true posterior's from reweighting, posterior induced by generated prior.

References

[1] A. Garbuno-Inigo, T. Helin, F. Hoffmann, and B. Hosseini. Bayesian posterior perturbation analysis with integral probability metrics. *arXiv preprint arXiv:2303.01512*, 2023.

A Dynamic Flux-Based Surrogate Approach to Interface Couplings Involving Full-Order Models and DMD-Based Surrogates

Edward Huynh[1], **Justin Owen**[2], **Paul Kuberry**[2], **and Pavel Bochev**[2]

[1] *The University of Texas at Austin, The Oden Institute for Computational Engineering and Sciences*
[2] *Sandia National Labs*

Partitioned methods for coupled problems rely on data transfers between subdomains to synchronize the subdomain equations and enable their independent solution. These methods enable code reuse, increase concurrency, and provide a convenient framework for plug-and-play multiphysics simulations. However, accuracy and stability of partitioned methods depends crucially on the type of information exchanged between the subproblems. Data-driven system identification methods allow the construction of accurate, computationally efficient and minimally intrusive data transfer surrogates. This approach leaves the main computational burden to an offline phase and consequently the application of this approach as the sole additional cost during the online simulation phase. In this presentation, we formulate and couple together fully data-driven (DMD-DMD) and hybrid (FOM-DMD) solvers for a model advection-diffusion transmission problem by using Dynamic Mode Decomposition (DMD) to learn the dynamics of both the subdomain states and the interface flux from data. Our results demonstrate that these couplings can be accomplished accurately at a cheaper computational cost.

Convergence Analysis for Infinite GMRES

J. Jackson[1] **and M. Embree**[1]

[1] *Department of Mathematics, Virginia Tech, Blacksburg, 24601, United States*

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 [1], converts these systems into infinite-dimensional systems with linear parameter dependence. This transformation involves a Carleman linearization-like 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, resulting from Carleman lifting, 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.

References

[1] E. Jarlebring and S. Correnty. Infinite GMRES for parametrized linear systems. *SIAM Matrix Anal. Appl.*, 43:1382–1405, 2022.

Toward a low-order representation of counter-rotating vortex pairs

J. Gregory Janes[1], Steffen Werner[2], Alistair D. G. Hales[1], and Justin W. Jaworski[1]

[1] *Virginia Tech, Department of Aerospace and Ocean Engineering, Blacksburg, VA 24061, USA.*

[2] *Virginia Tech, Department of Mathematics, Division of Computational Modeling and Data Analytics, and National Security Institute, Blacksburg, VA 24061, USA.*

The identification of dominant nonlinear dynamics in unequal nearly-parallel counter-rotating vortex filaments is an important problem in engineering. Dominant short and long-wavelength instabilities grow in coupled relationship, with small differences in initial perturbation and vortex parameters leading to vastly different vortex motion and deformations as the pair temporally evolves [1]. Traditional data-driven model reduction techniques, like exact dynamic mode decomposition (DMD), struggle at creating reduced order models (ROM) that capture the late-stage evolution of the vortex filaments. Furthermore, the computational strain imposed by the use of high spatial fidelity implicit large-eddy simulation (ILES) data makes it difficult to achieve high temporal resolution.

To address these challenges, we investigate the use of operator inference with roll outs, a technique introduced by Uy et al. [2] which resolves the sparsity and noise sensitivity of traditional methods. This approach defines a low-dimensional discrete flow map $\Phi_{\delta t}$ that advances the reduced state $\hat{q}$ through time:

$$\hat{q}_k(\mu) = \Phi_{\delta t}(\hat{q}_{k-1}(\mu), u_{k-1}(\mu), \mu; \theta). \tag{1}$$

The flow map $\Phi_{\delta t}$ is parameterized by operators θ that maintain a polynomial structure, separating linear from nonlinear dynamics to ensure model interpretability and scalability. Unlike static formulations that penalize residuals based only on a single step into the future, the roll out approach minimizes the discrepancy between projected observations $\bar{q}$ and predicted states $\hat{z}$ over a trajectory of length R:

$$\min_{\theta \in \Theta} J(\theta; \mu) = \sum_{k=0}^{K-R} \sum_{r=1}^{R} \|\bar{q}_{k+r}(\mu) - \hat{z}_{k,r}(\mu)\|_2^2. \tag{2}$$

The predicted states are generated by recursively applying the discrete flow map starting from an observed state:

$$\hat{z}_{k,r}(\mu) = \Phi_{\delta t}(\hat{z}_{k,r-1}(\mu), u_{k+r-1}(\mu), \mu; \theta), \quad \hat{z}_{k,0}(\mu) = \bar{q}_k(\mu). \tag{3}$$

For the counter-rotating vortex pair investigated here, this discrete-time formulation enables the learning of a robust quadratic model from sparsely sampled state observations. By leveraging automatic differentiation to optimize the non-convex objective, the learned operators accurately capture coupled instabilities without requiring the approximation of time derivatives from the high-fidelity snapshots.

References

[1] J. G. Janes, J. W. Jaworski, and A. D. G. Hales. On the evolution of unequal counter-rotating vortex pairs. In *AIAA AVIATION FORUM AND ASCEND 2025*, pages 1–13, Las Vegas, Nevada, July 2025. American Institute of Aeronautics and Astronautics.

[2] W. I. T. Uy, D. Hartmann, and B. Peherstorfer. Operator inference with roll outs for learning reduced models from scarce and low-quality data, Dec. 2022.

An Averaged Forward-Backward Dynamic Mode Decomposition for Noisy High-Dimensional Data

Aryeh Keating[1] **and Steffen W. R. Werner**[1, 2]

[1]*Department of Mathematics, Virginia Tech, Blacksburg, VA, USA*
[2]*Division of Computational Modeling and Data Analytics, and National Security Institute, Virginia Tech, Blacksburg, VA, USA*

The Dynamic Mode Decomposition (DMD) [2, 3] is a successful and popular approach for the inference of linear operators $\boldsymbol{A}$ from given data pairs $\{(\boldsymbol{x}_j, \boldsymbol{y}_j)\}_{j=1}^N$ so that

$$\boldsymbol{y}_j \approx \boldsymbol{A}\boldsymbol{x}_j, \quad \text{for } j = 1, \ldots, N.$$

However, the presence of measurement noise in the snapshot data $\{(\boldsymbol{x}_j, \boldsymbol{y}_j)\}_{j=1}^N$ quickly corrupts the accuracy of the results through statistical bias of the eigenvalues and modes of the inferred $\boldsymbol{A}$. There have been different approaches proposed to compensate for noise corruption, ranging from regularization of the underlying DMD least-squares problem to the truncation of corrupted data. One such method that is based on the linearity of the operator of interest is the Forward-Backward DMD (FBDMD) method. Thereby, the method computes the geometrical average of the forward operator $\boldsymbol{A}$ and the backward operator $\boldsymbol{B}$, for which $\boldsymbol{x}_j \approx \boldsymbol{B}\boldsymbol{y}_j$, allowing for the cancellation of the bias of the noise in the data; see [1]. However, FBDMD comes with its own challenges that result from the non-uniqueness of the geometrical average for matrices and the potentially small number of given data samples compared to the dimensions of the snapshots.

In this work, we propose an FBDMD operator based on the arithmetic average that achieves accuracy levels comparable to the classical FBDMD method, and we introduce a suitable projection framework that allows the computation of FBDMD for high-dimensional problems. The performance of the method is demonstrated in several numerical experiments in comparison to the classical DMD and FBDMD frameworks.

References

[1] S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley. Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *Exp. Fluids*, 57(3):42, 2016.

[2] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.*, 656:5–28, 2010.

[3] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: Theory and applications. *J. Comput. Dyn.*, 1(2):391–421, 2014.

Adaptive Reduced-Basis Trust-Region Methods for Defect Identification in Elastic Materials

Benedikt Klein[1] **and Mario Ohlberger**[1]

[1] *Applied Mathematics: Institute for Analysis and Numerics, University of Münster, Einsteinstr. 62, Münster, 48149Germany*

Monitoring the integrity of elastic structures using ultrasonic waves (Structural Health Monitoring, SHM) requires the efficient computation of mechanical parameters from measured surface displacements. The displacement field is governed by Cauchy's equation of motion, i.e., an elastic wave equation. Consequently, defect localization leads to a high-dimensional spatial parameter identification problem for a hyperbolic system with given initial and boundary conditions [4, 3]. Stable reconstructions typically rely on regularization techniques such as the iteratively regularized Gauss–Newton method (IRGNM). However, its practical deployment is computationally demanding due to the high-dimensional nature of the problem.
To address this bottleneck, we propose a reduced-order modeling approach that simultaneously reduces the state and parameter spaces using adaptively constructed reduced basis spaces. This yields online-efficient surrogate models for both the forward and adjoint evaluations required in derivative-based optimization. To ensure reliability, the IRGNM iteration is embedded into an adaptive, error-aware trust-region framework that certifies the accuracy of the reduced-order approximations. Our approach extends recent developments in [1, 2], which focus on elliptic and parabolic problems, to the hyperbolic setting. We demonstrate the reliability and effectiveness of the method for defect detection through numerical experiments.

References

[1] M. Kartmann, T. Keil, M. Ohlberger, S. Volkwein, and B. Kaltenbacher. Adaptive reduced basis trust region methods for parameter identification problems. *Computational Science and Engineering*, 1(1):3, 2024.

[2] M. Kartmann, B. Klein, M. Ohlberger, T. Schuster, and S. Volkwein. Adaptive reduced basis trust region methods for parabolic inverse problems. *arXiv preprint arXiv:2507.11130*, 2025.

[3] J. Seydel and T. Schuster. Identifying the stored energy of a hyperelastic structure by using an attenuated landweber method. *Inverse Problems*, 33(12):124004, 2017.

[4] J. Seydel and T. Schuster. On the linearization of identifying the stored energy function of a hyperelastic material from full knowledge of the displacement field. *Mathematical Methods in the Applied Sciences*, 40(1):183–204, 2017.

Streaming Operator Inference: Efficient Streaming Operator Learning for Large-Scale Dynamical Systems

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

[1] *Georgia Institute of Technology*
[2] *National Laboratory of the Rockies*

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) [1] 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 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. Numerical experiments on benchmark problems and 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 by over 99% and enabling dimension reductions exceeding 31,000×, resulting in orders-of-magnitude faster predictions. Our results establish Streaming OpInf as a scalable framework for reduced operator learning in large-scale and online streaming

References

[1] B. Peherstorfer and K. Willcox. Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*, 306:196–215, July 2016.

Stabilization Through Filtering for Non-Intrusive Reduced Order Models

Ian Moore[1] **and Traian Iliescu**[1]

[1] *Virginia Tech*

Operator inference (OpInf) is a non-intrusive reduced order modeling (ROM) framework that constructs dynamical system models directly from data by learning reduced operators, bypassing the need for intrusive access to full order model (FOM) codes [2]. While OpInf has demonstrated strong predictive capabilities across a range of applications, its performance can be sensitive to the choice of model hyper-parameters [1]. In practice, suboptimal tuning may lead to unstable dynamics or non-physical behavior in the learned reduced models. In this talk, we investigate the use of filtering strategies, drawing inspiration from ROM stabilization techniques and image processing methods to enhance the robustness of OpInf models. We explore how targeted filtering can mitigate instabilities and improve overall model fidelity without sacrificing the non-intrusive nature of the approach.

References

[1] I. Farcas, R. Munipalli, and K. E. Willcox. *On filtering in non-intrusive data-driven reduced-order modeling.*

[2] B. Peherstorfer and K. Willcox. Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*, 306:196–215, 2016.

Discovering Quadratic Representations for Nonlinear PDEs

A. Olivieri[1], G. Pogudin[2], and B. Kramer[1]

[1] *Department of Mechanical and Aerospace Engineering, University of California San Diego, CA, United States*
[2] *LIX, CNRS, École Polytechnique, Institute Polytechnique de Paris*

Nonpolynomial and nonquadratic PDEs are used to describe complex dynamical processes in science and engineering. Some examples are the cubic FitzHugh-Nagumo model, which explains the activation and deactivation dynamics of a spiking neuron; the cubic Brusselator model, used to predict oscillations in chemical reactions; the rational Euler equations that govern inviscid flows; and the nonadiabatic tubular reactor model with an exponential Arrhenius reaction term. Transforming or lifting nonquadratic models into quadratic ones has been used to simplify the study of such systems while allowing highly nonlinear dynamical behavior. In the context of model reduction, it was first introduced in [3] and has since been used to perform projection-based MOR [5], develop system-theoretic reduced models [1], and for structure-preserving [4] and data-driven ROMs [6]. There is a lack of computational tools for finding quadratic transformations for PDE systems; and so far, they have been derived by hand. This is tedious, error-prone, and often results in suboptimal lifted transformations.

Quadratization for PDEs is a process that transforms a nonquadratic PDE into a quadratic form by introducing auxiliary variables. The set of variables introduced is called a quadratization. For illustration, consider the PDE describing the evolution of the space and time-varying function $u(t, x)$:

$$u_t = u_x + u^3 \tag{1}$$

To quadratize (1) we introduce the variable $y(t, x) := u^2$. This allows us to write

$$y_t = 2u_t u = 2u_x u + 2u^4 = 2u_x u + 2y^2 \quad \text{and} \quad u_t = u_x + uy. \tag{2}$$

This is a quadratic equation in $u(t, x)$ and $y(t, x)$, so we call the set $\{u^2\}$ a quadratization for (1).

Multiple studies have focused on the ODE quadratization problem, and software tools are available to compute this transformation for any ODE [2]. However, no direct methods are available for the PDE case. We present theoretical results for the PDE quadratization problem, along with an algorithm and accompanying software that finds quadratizations of polynomial and rational spatially one-dimensional PDEs. The presented algorithm searches a combinatorial tree of possible transformations, uses branch-and-bound techniques to curb its computational complexity, and outputs a relative small set of variables that effectively quadratize a PDE system. To the best of our knowledge, these are the first results of automated quadratization for PDEs.

References

[1] P. Benner and T. Breiten. Two-sided projection methods for nonlinear model order reduction. *SIAM Journal on Scientific Computing*, 37(2):B239–B260, 2015.

[2] A. Bychkov and G. Pogudin. Optimal monomial quadratization for ODE systems. In *Combinatorial Algorithms*, pages 122–136, 2021.

[3] C. Gu. QLMOR: A projection-based nonlinear model order reduction approach using quadratic-linear representation of nonlinear systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(9):1307–1320, 2011.

[4] S. Harsh, J. D. Draxl Giannoni, and B. Kramer. Nonlinear energy-preserving model reduction with lifting transformations that quadratize the energy. *Physica D*, 483:134954, 2025.

[5] B. Kramer and K. E. Willcox. Nonlinear model order reduction via lifting transformations and proper orthogonal decomposition. *AIAA Journal*, 57(6):2297–2307, 2019.

[6] E. Qian, B. Kramer, B. Peherstorfer, and K. Willcox. Lift & learn: Physics-informed machine learning for large-scale nonlinear dynamical systems. *Physica D*, 406:132401, 2020.

Matrix Decomposition for Characterizing Complex Vortex Structures

A. Palmiero[1], J. Takei[2], M. Al Shaaban[3], L. Dereje[4], S. Panitch[5], C. Chen[6], and Y. Yan[7]

[1]*Department of Mathematics and Statistics, University of Maryland Baltimore County*
[2,4,5]*Department of Mathematics, Yale University*
[3]*Department of Mathematics, Brandeis University*
[6,7]*Department of Mathematics and Statistical Sciences, Jackson State University*

The accurate characterization of complex spatiotemporal structures, such as coherent vortices in high-speed turbulent boundary layers, poses a significant challenge due to the high dimensionality and nonlinear nature of the governing Navier-Stokes equations. To circumvent these issues, flows are computer-simulated over a grid, yielding discrete data which captures physical characteristics of the flow at each grid-point. Through the decomposition of select attributes of this physical data, lower-dimensional matrices of temporal coefficients and spatial modes can be derived. Our work focuses on optimizing matrix decomposition such that the resulting spatial modes best capture dominant physical structures within the fluid flow. The main issue encountered when using a dimensionality reduction technique in this manner is the appearance of pseudo-vortices, a truncation error from the unresolved high-frequency spatial modes.

In this study, we explore various matrix decomposition techniques to reconstruct complex vortex structures while minimizing pseudo-vortices. We extend the use of Non-negative Matrix Factorization (NMF) into the domain of Fluid Dynamics and compare the performance of NMF to Proper Orthogonal Decomposition (POD) in accurately reconstructing the fluid flow. NMF visually eliminates pseudo-vortices better than POD when reconstructing fluid flows, however spatial modes generated using NMF have a tendency to exactly replicate physical features of the original fluid flow, rather than capturing abstract components which are not seen in the original vortex structure, like POD. Our methodology establishes a computationally efficient platform for investigating vortex dynamics, offering substantial reductions in data dimensionality while preserving essential features of turbulent transport.

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

Heike Faßbender[1], Serkan Gugercin[2], and Till Peters[1]

[1]*Institute for Numerical Analysis, TU Braunschweig*
[2]*Department of Mathematics, Virginia Tech*

Today, mathematical modeling is dominated by increasingly high-dimensional and complex dynamical systems. One special type of structure is the bilinear state equation which naturally appears in various applications or as a result of a Carleman bilinearization of nonlinear dynamics; see, e.g., [7]. 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; see, e.g., [6]. As a combination, we study the so-called bilinear quadratic output (BQO) systems described as

$$\dot{x}(t) = Ax(t) + \sum_{k=1}^{m} N_k x(t) u_k(t) + Bu(t), \tag{1a}$$

$$y(t) = Cx(t) + \begin{bmatrix} x(t)^T M_1 x(t) \\ \vdots \\ x(t)^T M_p x(t) \end{bmatrix} \tag{1b}$$

with state $x(t) \in \mathbb{R}^n$, input $u(t) \in \mathbb{R}^m$, output $y(t) \in \mathbb{R}^p$ and matrices $A \in \mathbb{R}^{n\times n}, B \in \mathbb{R}^{n\times m}, C \in \mathbb{R}^{p\times n}, N_k \in \mathbb{R}^{n\times n}$ for $k = 1, \ldots, m$, $M_j \in \mathbb{R}^{n\times n}$ for $j = 1, \ldots, p$, and $t \in T := [0, \infty)$.
BQO systems are a generalization of both bilinear systems [1, 8] and linear quadratic output (LQO) systems [3]. More specifically, in (1b), setting $M_1 = \cdots = M_p = 0$ leads to the linear output $y(t) = Cx(t)$, recovering the bilinear dynamical systems. Similarly, setting $N_1 = \cdots = N_m = 0$ in (1a) leads to the linear state equation $\dot{x}(t) = Ax(t) + Bu(t)$, recovering the linear quadratic output (LQO) systems. For large dimensional systems, the task of structure-preserving model order reduction arises. Thus, we seek a high-fidelity BQO system of reduced dimension that approximates the original input-output behavior well. Whereas a balancing approach has been already discussed in [4], here we focus on an $\mathcal{H}_2$ optimal approach, similar to $\mathcal{H}_2$ optimal reduction for bilinear systems; see, e.g., [8, 2, 5]. In this talk, 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] S. Al-Baiyat and M. Bettayeb. A new model reduction scheme for k-power bilinear systems. In *Proceedings of 32nd IEEE Conference on Decision and Control*, pages 22–27 vol.1, Dec. 1993.

[2] P. Benner and T. Breiten. Interpolation-based $\mathcal{H}_2$-model reduction of bilinear control systems. *SIAM Journal on Matrix Analysis and Applications*, 33(3):859–885, 2012.

[3] P. Benner, P. Goyal, and I. P. Duff. Gramians, Energy Functionals, and Balanced Truncation for Linear Dynamical Systems With Quadratic Outputs. *IEEE Transactions on Automatic Control*, 67(2):886–893, Feb. 2022. Conference Name: IEEE Transactions on Automatic Control.

[4] H. Faßbender, S. Gugercin, and T. Peters. Bilinear Quadratic Output Systems and Balanced Truncation, July 2025. arXiv:2507.03684 [math].

[5] G. Flagg and S. Gugercin. Multipoint volterra series interpolation and $\mathcal{H}\_2$ optimal model reduction of bilinear systems. *SIAM Journal on Matrix Analysis and Applications*, 36(2):549–579, 2015.

[6] R. Pulch and A. Narayan. Balanced truncation for model order reduction of linear dynamical systems with quadratic outputs. *SIAM Journal on Scientific Computing*, 41(4):A2270–A2295, 2019.

[7] W. J. Rugh. *Nonlinear System Theory: The Volterra/Wiener Approach.* Johns Hopkins University Press, 1981.

[8] L. Zhang and J. Lam. On $H2$ model reduction of bilinear systems. *Automatica*, 38(2):205–216, Feb. 2002.

Towards a Black-Box MPS Laplacian Eigensolver

Hayden Ringer[1], Mark Embree[1], and Justin Krometis[2]

[1] *Virginia Tech, Department of Mathematics*
[2] *Virginia Tech, Hume Center for National Security and Technology*

Consider the standard Dirichlet Laplace eigenproblem for a plane domain Ω with Lipschitz boundary:

$$-\Delta u = \lambda u \quad \text{on } \Omega \tag{1}$$

$$u = 0 \quad \text{on } \partial\Omega \tag{2}$$

Denote the eigenvalues corresponding to the Dirichlet Laplacian on Ω as $\mathrm{Spec}(-\Delta_\Omega^D) = \{\lambda_1, \lambda_2, \dots\}$.

A *spectral optimization problem* or *spectral inverse problem* can generally be written in the form:

$$\min_{\Omega \in \mathcal{C}} L\big(\mathrm{Spec}(-\Delta_\Omega^D)\big)\,, \tag{3}$$

where L is a suitable loss function and $\mathcal{C}$ is a class of domains. Numerical solutions to spectral optimization or spectral inverse problems necessarily depend on reliable forward eigensolvers. The Method of Particular Solutions (MPS), introduced by in [2] and "revived" by Betcke and Trefethen in [1], is a powerful mesh-free technique for accurate approximation of Dirichlet Laplacian eigenvalues for planar domains. However, it is not easily made into a "black-box" eigensolver, and in fact requires significant domain-specific parameter tuning.

In this talk, we discuss our work in developing an MPS implementation which can reliably compute a desired number of eigenvalues of a given domain Ω to user-specified precision, without requiring the user to manually specify MPS parameters. We show that, for most domains, it is feasible to obtain 8 digits of floating point accuracy in the "low spectrum" using our approach.

References

[1] T. Betcke and L. N. Trefethen. Reviving the Method of Particular Solutions. 47(3):469–491.

[2] L. Fox, P. Henrici, and C. Moler. Approximations and Bounds for Eigenvalues of Elliptic Operators. 4(1):89–102.

Likelihood-informed model reduction for linear Bayesian inverse problems

J. Scheffels[1], E. Qian[2], I. Papaioannou[1], and E. Ullmann[3]

[1] *Engineering Risk Analysis Group, TU Munich*
[2] *School of Aerospace Engineering and School of Computational Science and Engineering, Georgia Institute of Technology*
[3] *Department of Mathematics, TU Munich*

The Bayesian approach to solving inverse problems updates a prior distribution to a posterior distribution of the parameters through data [1]. The likelihood of the data needed to obtain the posterior distribution depends on a computational model for the numerical solution of a partial differential equation. Standard procedures for solving the Bayesian inverse problem are based on multiple evaluations of the likelihood function and, hence, the underlying computational model. The application of model order reduction enables a cheaper solution of the inverse problem for resource demanding computational models.

We focus on Bayesian linear inverse problems with Gaussian parameter priors and sparse data. As limited data is only informative in a low-dimensional subspace of the parameter space, the posterior in such problems can be expressed as a low-rank update of the prior distribution [3]. In this work, we introduce a framework that leverages the low-rank mapping between the data and the parameter itself to construct a reduced-order model for inverse problems [2]. This approach is based on the leading eigendirections of the Fisher information matrix and the prior precision and is referred to as likelihood-informed subspace (LIS). Through numerical examples drawn from structural engineering, we compare the LIS reduced-order model with reduced-order models obtained using proper orthogonal decomposition (POD) and present a-priori error bounds for the posterior approximation.

References

[1] M. Dashti and A. M. Stuart. *The Bayesian Approach to Inverse Problems*. Springer International Publishing, 2017.

[2] J. Scheffels et al. Likelihood-informed Model Reduction for Bayesian Inference of Static Structural Loads. *Arxiv*, 2025.

[3] A. Spantini et al. Optimal Low-Rank Approximations of Bayesian Linear Inverse Problems. *SIAM Journal on Scientific Computing*, 37:A2451–A2487, 2015.

Localized Model Order Reduction for Symmetric Generalized Eigenvalue Problems: A Neutron Diffusion Case Study

Francesco A.B. Silva[1]

[1] *Department of Nuclear Engineering, Texas A&M University, 423 Spence St, College Station, 77843-3133, TX, USA*

In this work, we develop a localized model order reduction framework for symmetric generalized eigenvalue problems arising from the discretization of neutron diffusion equations on modular domains, in the spirit of static condensation and generalized multiscale finite element methods [1, 2, 3]. As in these approaches, we employ a domain decomposition formulation in which global bilinear forms are expressed as sums of local contributions; restriction and continuation operators connect local and global variables, and the dominant eigenpair is characterized via a global Rayleigh quotient. Two families of local eigenvalue problems, with Neumann or Dirichlet conditions on internal interfaces, provide rigorous upper and lower bounds for the global dominant eigenvalue.
On each modular subdomain, we define outer-to-inner transfer operators that map data from oversampled boundaries to internal interfaces. Their spectral properties are used to construct Kolmogorov-optimal reduced boundary spaces and associated localized basis functions for tentative training values of the global dominant eigenvalue that lie within a certified admissible interval. A spectral greedy strategy then selects parameter-robust local approximation spaces on each subdomain and assembles a globally reduced generalized eigenvalue problem by projecting the original formulation onto the span of the selected bases.
The methodology is demonstrated for the one-group neutron diffusion eigenvalue problem in a heterogeneous, modular, reactor-like configuration. Numerical experiments show that the resulting reduced spaces retain full-order accuracy for the dominant eigenpair while requiring only a small number of localized basis functions per interface.

References

[1] A. Buhr, L. Iapichino, M. Ohlberger, S. Rave, F. Schindler, and K. Smetana. *6 Localized model reduction for parameterized problems*, pages 245–306. De Gruyter, Berlin, Boston, 2021.

[2] Y. Efendiev, J. Galvis, and T. Y. Hou. Generalized multiscale finite element methods (gmsfem). *Journal of Computational Physics*, 251:116–135, 2013.

[3] D. B. Phuong Huynh, D. J. Knezevic, and A. T. Patera. A Static condensation Reduced Basis Element method: approximation and *a posteriori* error estimation. *ESAIM: Mathematical Modelling and Numerical Analysis*, 47(1):213–251, 2013.

A Tangential ADI Method for the Efficient Solution of Large-scale Indefinite Lyapunov Equations

Rudi Smith[1] **and Steffen W. R. Werner**[1,2]

[1]*Department of Mathematics, Virginia Tech, Blacksburg, VA, USA*
[2]*Division of Computational Modeling and Data Analytics, and National Security Institute, Virginia Tech, Blacksburg, VA, USA*

Introduction. Continuous-time algebraic Lyapunov equations are foundational tools in systems and control theory with various applications in model order reduction [3] and controller design [1]. The Low-Rank Alternating Direction Implicit (LR-ADI) method [2, 4] is an effective approach to solve large-scale Lyapunov equations. We focus on the generalized case with an indefinite right-hand side:

$$AXE^{\mathsf{T}} + EXA^{\mathsf{T}} = -BRB^{\mathsf{T}}, \tag{1}$$

where $A, E \in \mathbb{R}^{n \times n}$ are sparse, $B \in \mathbb{R}^{n \times m}$, and $R \in \mathbb{R}^{m \times m}$ is symmetric indefinite. In the case of medium and high-rank right-hand sides, standard ADI naturally suffers from increasing computational costs due to the block-wise growth of the constructed solution factors ($X \approx LDL^{\mathsf{T}}$) by m columns at every step.

Tangential Framework. We present a tangential ADI framework for indefinite Lyapunov equations that overcomes this limitation by replacing block updates with a sequence of memory-efficient rank-1 updates. By mitigating the rapid growth of the solution dimensions, our approach offers a more scalable algorithm.

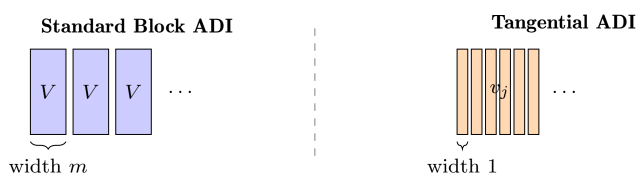

At each step, an adaptive shift α_j and a corresponding tangential direction t_j are selected to column-wise construct the solution. The implicit residual formulation avoids forming the dense full-dimensional solution matrix, making the approach usable for large-scale sparse coefficients.

Numerical Results. We demonstrate the performance of the proposed method using well-known benchmark problems from the literature. Experiments confirm that the tangential approach matches the residual convergence of block ADI while producing solution factors with significantly lower rank, validating the method for practical applications.

References

[1] E. Armstrong. An extension of Bass' algorithm for stabilizing linear continuous constant systems. *IEEE Trans. Autom. Control*, 20(1):153–154, 1975.

[2] P. Benner, J.-R. Li, and T. Penzl. Numerical solution of large-scale lyapunov equations, riccati equations, and linear-quadratic optimal control problems. *Numerical Linear Algebra with Applications*, 15(9):755–777, 2008.

[3] P. Benner, V. Mehrmann, and D. C. Sorensen. *Dimension Reduction of Large-Scale Systems*, volume 45 of *Lect. Notes Comput. Sci. Eng.* Springer, Berlin, Heidelberg, 2005.

[4] N. Lang, H. Mena, and J. Saak. An LDL^T factorization based ADI algorithm for solving large-scale differential matrix equations. *Proc. Appl. Math. Mech.*, 14(1):827–828, 2014.

Wavelet-Based Observables for Koopman Analysis

C. Tilki[1] **and S. Güğercin**[1]

[1] *Department of Mathematics, Virginia Tech, Blacksburg, 24061, VA, United States*

Abstract

Dynamic Mode Decomposition (DMD) [2] is a widely used data driven modeling technique for approximating linear representations of nonlinear dynamical systems of the form

$$\dot{\boldsymbol{x}} = \mathcal{T}(\boldsymbol{x}), \qquad \boldsymbol{x} \in \mathcal{M} \subseteq \mathbb{R}^n.$$

Its connection to Koopman operator theory has led to the Extended DMD (EDMD) method [3], which approximates the action of the Koopman operator by a Galerkin projection onto a finite-dimensional subspace spanned by a prescribed set of observables $\{\psi_j\}_{j=1}^M$, that are scalar valued functions of the form $\psi_j : \mathcal{M} \to \mathbb{C}$.

Given snapshot data $\{\boldsymbol{x}(t_i)\}_{i=0}^N \subseteq \mathcal{M}$, EDMD computes a finite-dimensional approximation of the (projected) Koopman operator by solving the least-squares problem

$$\boldsymbol{K} = \underset{\hat{\boldsymbol{K}} \in \mathbb{C}^{M \times M}}{\arg\min} \, \|\Psi_1 - \hat{\boldsymbol{K}} \Psi_0\|_F^2, \text{ where } \Psi_j := \begin{bmatrix} \psi_1(\boldsymbol{x}(t_j)) & \psi_1(\boldsymbol{x}(t_{j+1})) & \dots & \psi_1(\boldsymbol{x}(t_{j+N-1})) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_M(\boldsymbol{x}(t_j)) & \psi_M(\boldsymbol{x}(t_{j+1})) & \dots & \psi_M(\boldsymbol{x}(t_{j+N-1})) \end{bmatrix}.$$

The resulting matrix $\boldsymbol{K}$ serves as a finite-dimensional approximation of the projected Koopman operator.

Accuracy of this approximation depends critically on the choice of observables and the approximation error can be decomposed into projection and sampling errors. In particular, commonly used observable sets may introduce spurious data-independent errors beyond the intrinsic Galerkin projection. Hence in practice, selecting effective observables ψ remains a major challenge.

In this work, we construct a class of *wavelet-based observables* for EDMD, inspired by the lifting mechanism underlying the Wavelet-based DMD (WDMD) framework [1]. We show that the resulting approximation space is Koopman invariant, ensuring that EDMD recovers the exact action of the Koopman operator on this subspace in the infinite-data limit. Consequently, the proposed approach avoids spurious data-independent errors and yields a consistent Galerkin restriction of the Koopman operator. Since wavelet-based observables eliminate these data-independent errors, they become a good candidate for EDMD observables.

Building on this, we introduce the wavelet-based DMD via continuous wavelet transform (cWDMD) algorithm which employs wavelet-based observables within EDMD and investigate its numerical performance on the chaotic Lorenz 67 system. The results demonstrate that the proposed observables lead to improved Koopman operator approximations in regimes characterized by strong nonlinearity and chaotic dynamics, in agreement with the theoretical analysis.

References

[1] M. Krishnan, S. Güğercin, and P. A. Tarazaga. A wavelet-based dynamic mode decomposition for modeling mechanical systems from partial observations. *Mechanical Systems and Signal Processing*, 187:109919, 2023.

[2] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.

[3] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

Model Order Reduction using Group Convolutional Autoencoders

Sven Ullmann[1,2], Mario Ohlberger[2], and Stephan Rave[2]

[1] *Institute of Applied Analysis and Numerical Simulation, University of Stuttgart, Pfaffenwaldring 57, Stuttgart 70569, Germany*

[2] *Institute for Analysis and Numerics and Mathematics Münster, University of Münster, Einsteinstrasse 62, 48149 Münster, Germany*

MOR is an important tool for the efficient simulation of parametrized and dynamical systems governed by PDEs. Classical projection-based MOR techniques rely on linear trial spaces, typically constructed via methods such as POD. While highly effective for diffusion-dominated problems, linear subspace approaches suffer from a fundamental expressivity limitation when applied to advection-dominated problems, which are characterized by a slowly decaying Kolmogorov n-width. To overcome these limitations, nonlinear MOR techniques have gained significant attention. A prominent realization of this idea is the use of convolutional autoencoders to learn nonlinear reduced manifolds directly from snapshot data, leading to a manifold Galerkin formulation [3].

In these convolutional networks layers are by definition translation equivariant, meaning if the input image of a layer is shifted and then convoluted, this results in the same as first convoluting and then shifting. Interpreting translations as the action of the group $G = (\mathbb{R}^2, +)$ acting on $\mathbb{R}^2$, a convolutional layer Φ is an equivariant map, i.e., $\Phi(T_y f) = T_y \Phi(f)$ for all $y \in G$., where $(T_y f)(x) := f(x - y)$ denotes the action of the translation group. This concept has been generalized through the introduction of group convolutions [2], resulting in layers that are additionally equivariant w.r.t to symmetry groups such as rotations or reflections. Originally developed in the context of classification, group equivariant networks provide a systematic way to encode symmetry directly into the network architecture.

In this work, we investigate group equivariant autoencoders for the construction of nonlinear trial manifolds. We demonstrate the approach on a 2D linear wave equation describing a thin pulse, parametrized by the wave speed - a problem known to exhibit a slowly decaying Kolmogorov n-width and thus to require nonlinear reduction techniques. By exploiting rotational equivariance in the encoder and decoder architectures, the learned reduced manifold generalizes beyond the training data and enables accurate ROMs for rotated wave configurations without retraining. Since the linear wave equation admits a Hamiltonian formulation, we further incorporate the ideas of [1] by enforcing weak symplecticity of the decoder. This induces a symplectic structure on the trial manifold and results in ROMs that, for example, conserve energy.

Literatur

[1] P. Buchfink, S. Glas, and B. Haasdonk. Symplectic model reduction of hamiltonian systems on nonlinear manifolds and approximation with weakly symplectic autoencoder. *SIAM Journal on Scientific Computing*, 45(2):A289–A311, 2023.

[2] T. Cohen and M. Welling. Group equivariant convolutional networks. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999, New York, New York, USA, 20–22 Jun 2016. PMLR.

[3] K. Lee and K. T. Carlberg. Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics*, 404:108973, 2020.

Tensor-based multivariate rational approximation

Leonie Van Pottelberghe, Daan Huybrechs and Karl Meerbergen[1]

[1] *Dept. of Computer Science, KU Leuven, 3001 Leuven, Belgium*

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.

Inference-Oriented Nonlinear Model Reduction for Bayesian Data Assimilation

Weiting Yi[1] **and Elizabeth Qian**[1]

[1] *Georgia Institute of Technology, GA, USA*

A central class of data assimilation problems concerns inferring the initial conditions of a dynamical system from noisy measurements collected at different time steps. Bayesian inference formulations are attractive because they yield posterior distributions that quantify the uncertainty in the estimated initial conditions given the measurements. However, in nonlinear dynamical systems, these posterior distributions are usually not available in closed form and require computationally intensive sampling procedures such as MCMC, which involve repeated simulations of the dynamical system for many candidate initial conditions. Such procedures become prohibitively expensive in high-dimensional nonlinear dynamical systems.

Successful Bayesian approaches for large-scale nonlinear data assimilation problems therefore often combine sampling-based inference tools with surrogate models for high-dimensional dynamical systems. Projection-based model reduction constructs low-dimensional surrogate models by projecting system operators onto a *low-dimensional subspace*, yielding reduced systems that are cheap to simulate and significantly accelerate posterior sampling, but such techniques (e.g., principal component analysis) optimize state-space accuracy rather than posterior fidelity. An alternative strategy reduces the dimensionality of the inference problem by identifying a *likelihood-informed subspace* (LIS) spanned by parameter directions most informed by the measurements, enabling more efficient posterior sampling in this low-dimensional subspace [1]. The LIS approach yields provably optimal posterior approximations in the linear Gaussian initial-condition inference setting but fails to reduce simulation cost in sampling-based Bayesian inference tools.

In this work, we propose a projection-based nonlinear model reduction framework for Bayesian initial-condition inference by exploiting the LIS. Specifically, we project the nonlinear dynamical system onto the LIS to obtain a reduced system that preserves the most measurement-informed initial-condition directions and lowers the cost of forward simulations during sampling. This framework provides two complementary benefits: it reduces the computational cost of posterior sampling through the reduced-order system and improves posterior approximation accuracy by aligning the model reduction with measurement-informed directions of the initial conditions. The linear instance of this framework was studied in [2], where numerical and theoretical results demonstrated that the resulting inference-oriented linear model reduction achieves near-optimal posterior covariance approximations. Experiments on benchmark nonlinear model reduction and data assimilation problems demonstrate that our framework reduces computational cost in posterior sampling while yielding more accurate posterior approximations than general projection-based model reduction techniques.

References

[1] T. Cui, J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015, 2014.

[2] E. Qian, J. M. Tabeart, C. Beattie, S. Gugercin, J. Jiang, P. R. Kramer, and A. Narayan. Model reduction of linear dynamical systems via balancing for bayesian inference. *Journal of Scientific Computing*, 91(1):29, 2022.