siam 2026

Searchable Abstracts Document

SIAM Conference on Uncertainty Quantification (UQ26)

March 22–25, 2026

Hyatt Regency Minneapolis
Minneapolis, Minnesota, U.S.

This document was current as of March 11, 2026. Abstracts appear as submitted.

siam

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP1

Denoising Diffusion Models: From Generative Modelling to Bayesian Computation

Recent breakthroughs in generative AI have demonstrated a remarkable ability to produce photorealistic images and videos, coherent music, and even novel protein structures. At the heart of many of these advances is a powerful and elegant framework: diffusion models. This talk will provide an overview of the core "denoising" mechanism that powers these models. We will then move beyond their well-known generative capabilities to explore their flexibility as a powerful computational tool. We will demonstrate how this framework can be leveraged to solve challenging inverse problems, discuss its deep connections to computational optimal transport, and examine how related ideas are being successfully explored to sample from complex, high-dimensional distributions. Finally, the talk will address the current limitations of these models and discuss some key open challenges.

Arnaud Doucet
Google DeepMind, United Kingdom
doucet@stats.ox.ac.uk

IP2

Non-intrusive Structural-Preserving Sequential Data Assimilation

Data assimilation (DA) methods combine model predictions with observational data to improve state estimation in dynamical systems, inspiring their increasingly prominent role in geophysical and climate applications. Classical DA methods assume that the governing equations modeling the dynamics are known, which may not be true in real world applications. Machine learning (ML) provides a flexible alternative by learning surrogate models directly from data, but standard ML methods struggle in noisy and data-scarce environments, where meaningful extrapolation requires incorporating physical constraints. Recent advances in structure-preserving ML architectures, such as the development of the entropy-stable conservative flux form network (ESCFN), highlight the critical role of physical structure in improving learning stability and accuracy for unknown systems of conservation laws. Structural information has also been shown to improve DA performance. Gradient-based measures of spatial variability, in particular, can help refine ensemble updates in discontinuous systems. Motivated by both of these recent innovations, we propose a new non-intrusive, structure-preserving sequential data assimilation (NSSDA) framework that leverages structure at both the forecast and analysis stages. Our method operates in a highly constrained environment, using only a single noisy trajectory for both training and assimilation.

Anne Gelb
Dartmouth College
Department of Mathematics
Anne.E.Gelb@Dartmouth.edu

IP3

Optimal Sampling for Linear and Nonlinear Approximation

This talk explores the approximation of functions from limited information, using linear or nonlinear tools. For linear approximation in L^2 from point evaluations, recent results show that weighted least-squares projections deliver quasi-optimal approximations with near to optimal sampling budgets, using i.i.d. samples from distributions tailored to the linear approximation tool, as well as subsampling methods. The first part of this talk surveys recent advances in optimal sampling methods for linear approximation, extending beyond the L^2 setting to encompass more general types of information and related Hilbert spaces—including those relevant to physics-informed machine learning—and going beyond i.i.d. sampling. In the second part, we show how these methods can be used to approximate functions with nonlinear approximation tools with differentiable parameterizations, including neural networks or tensor networks. This involves leveraging optimization algorithms on manifolds and optimal sampling methods for the projection of gradients onto successive linear spaces. Finally, we address the challenging problem of sampling from optimal distributions, particularly in contexts where only generating systems of linear spaces are accessible—a scenario commonly encountered when approximating with neural networks or more general compositional functions.

Anthony Nouy
Centrale Nantes, Nantes Université, France
anthony.nouy@ec-nantes.fr

IP4

NUTS for You: New Advances in No-U-Turn Samplers and the Future of Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) remains a central tool for uncertainty quantification, yet the efficiency of classical algorithms often deteriorates in high dimensions or anisotropic geometries. The No-U-Turn Sampler (NUTS) and its descendants have reshaped practical Bayesian computation by learning local geometry on the fly, enabling efficient exploration even in complex, high-dimensional landscapes. This talk revisits the mathematical foundations of NUTS and shows how they can be extended and unified within a broader adaptive framework. This perspective leads to new algorithms that preserve the self-tuning spirit of NUTS while extending its reach to curved and discrete geometries. I will share recent insights into why these methods mix so efficiently, drawing on ideas from geometry, probability, and analysis. Along the way, we will see how No-U-Turn ideas are evolving from clever computational innovations into a principled theory of locally adaptive MCMC, bringing us closer to the long-standing goal of samplers that require no tuning and are provably efficient.

Nawaf Bou-Rabee
Rutgers UniversityCamden and Flatiron Institute, U.S.
nb361@camden.rutgers.edu

IP5

CANCELLED: The Role of Gaussian Markov Random Fields in Statistical Modeling and Scalable Bayesian Inference

Gaussian Markov random fields (GMRFs) play a fundamental role in modern uncertainty quantification and Bayesian inference. In this talk, I will discuss their critical importance in statistical modeling, latent Gaussian models, and Gaussian processes, where GMRFs enable computational efficiency through sparse linear algebra. A GMRF is a Gaussian distributed vector x with a Markovian structure, where selected elements x_i and x_j are conditionally independent given the remaining elements. The Markov property naturally arises through a constructive formula-

tion of correlation or from Gaussian fields that are defined via solutions to stochastic partial differential equations using finite element methods. GMRFs induces sparsity in the associated symmetric positive-definite precision matrices, a property that can be exploited in the necessary key operations of factorization and selected matrix inversion, enabling scalable inference for high-dimensional problems. Thus making GMRFs a cornerstone in Bayesian inference, where computational cost is a central limitation.

Håvard Rue
King Abdullah University of Science and Technology
haavard.rue@kaust.edu.sa

IP6

30 Years of Uncertainty Quantification: The Evolution of UQ Methods, Example Applications, and Future Challenges

This talk will present the evolution of uncertainty quantification (UQ) and sensitivity analysis (SA) methods over the past thirty years, provide some example applications, and outline future challenges. A brief history of UQ/SA in the computational science community will be discussed, along with development of methods such as stochastic expansions, multifidelity approaches, and variance-based decomposition. Example applications of UQ and SA will be presented, followed by an outline of challenges that still remain for the practitioner when performing UQ/SA studies, such as in the context of machine learning algorithms or digital twins.

Laura P. Swiler
Retired (Sandia National Laboratories)
lpswiler@msn.com

IP7

Quasi-Monte Carlo for Precision Oncology

Quasi-Monte Carlo (QMC) methods have been successful for high dimensional integration, multivariate function approximation, density estimation, etc., in many application areas including uncertainty quantification problems driven by partial differential equations (PDEs) with random coefficients. This talk will showcase our latest work on the application of "lattice" QMC methods to a class of semilinear parabolic reaction-diffusion PDEs used to model tumor growth and treatment. Mathematical models of tumor growth are largely phenomenological in nature, capturing infiltration of the tumor into surrounding healthy tissue, proliferation of the existing tumor, and patient response to therapies, such as chemotherapy and radiotherapy. Considerable inter-patient variability, inherent heterogeneity of the disease, sparse and noisy data collection, and model inadequacy all contribute to significant uncertainty in the model parameters. It is crucial that these uncertainties can be efficiently propagated through the model to compute quantities of interest (QoIs), which in turn may be used to inform clinical decisions. We will demonstrate both theoretically and computationally the success of QMC for this challenging nonlinear PDE problem. This talk is based on joint work with Alexander Gilbert (UNSW Sydney), Dirk Nuyens (KU Leuven), Graham Pash (UT Austin), Ian Sloan (UNSW Sydney), and Karen Willcox (UT Austin).

Frances Y. Kuo
UNSW Sydney, Australia
f.kuo@unsw.edu.au

IP8

The Many Lives of Gaussian Processes

Beyond their foundational roots in statistical and stochastic process theory, Gaussian processes (GPs) are now central to a versatile and widely used framework spanning applied mathematics and probabilistic machine learning. Despite well-known limitations, GP-based methods continue to show remarkable adaptability, often outlasting shifting methodological trends. In this talk, I will highlight developments in modeling, prediction, and active learning with GPs, offering perspectives on their enduring role as a powerful building block for function approximation and uncertainty quantification.

David Ginsbourger
University of Bern
david.ginsbourger@unibe.ch

SP0

Outreach Talk - Mathematical Sciences at NSF: Strategic Directions and NSF Funding Updates

This presentation outlines the primary goals and recent developments of the Mathematical Sciences Sections (MSS) at the NSF, with particular emphasis on the Computational Mathematics program. The talk highlights the role of MSS in supporting foundational research and their connections to modern uncertainty quantification efforts. An up-to-date roadmap of NSF funding opportunities relevant to the computational mathematics community will also be presented. The session aims to help the community better align their research goals with national funding priorities.

Yuliya Gorb
National Science Foundation
ygorb@nsf.gov

SP1

SIAG/Uncertainty Quantification Early Career Prize Lecture - Surrogate Modeling for Stochastic Simulators: A Trajectory-based Spectral Approach

Many computer models involve sources of uncertainty that cannot be meaningfully represented through uncertain input parameters alone. Examples include simulations driven by environmental variability, such as wind or earthquakes, as well as agent-based models whose randomness arises from internal decision processes. Such stochastic simulators produce inherently random responses, posing fundamental challenges to standard uncertainty quantification and surrogate modeling techniques. While surrogate modeling for deterministic simulators is well established, corresponding methods for stochastic simulators are far less mature. In this talk, we first provide a high-level overview of existing approaches and the different modeling paradigms they employ. We then introduce the key ideas behind our recently developed trajectory-based spectral approach, which treats entire stochastic responses as the object of approximation rather than modeling output distributions pointwise in parameter space. Finally, we show how interactions between inputs and stochastic effects reveal key structure in the simulator's behavior and naturally govern the complexity of the surrogate model. Overall, this trajectory-based framework produces surrogates for stochastic simulators that capture both paramet-

ric dependence and intrinsic stochastic behavior in an interpretable manner.

Nora Lüthen
ETH Zürich
luethen@ibk.baug.ethz.ch

CP1

Adversarial Calibration of Ipcw-Weighted Bayesian Neural Mixtures for Right-Censored Survival Prediction

Right-censored survival data are ubiquitous in clinical studies and pose two intertwined challenges: censoring biasskewing risk estimates when events go unobservedand model overconfidencewhere deep survival networks underestimate uncertainty. To tackle both, we propose an adversarial calibration framework that (i) applies inverse-probability-of-censoring weights to correct partial-likelihood bias, (ii) replaces the standard DeepSurv output with a Bayesian mixture of parametric hazard components to capture parameter and structural uncertainty, and (iii) trains a generative uncertainty critic to distinguish true event-time samples from model-drawn ones and thereby correct overconfident predictions. We evaluate on two large clinical cohorts (METABRIC breast cancer and SEER-Medicare lung cancer), comparing against standard DeepSurv+IPCW, DeepSurv+Bayesian mixtures (no adversary), and common post-hoc calibration methods (temperature scaling, isotonic regression). Our approach achieves up to a 15% reduction in integrated Brier score, calibration slopes closer to 1.0, and 2–3-point gains in concordance indexespecially in low-incidence subgroups (¡10% event rate). Ablation studies confirm the critic uncovers latent miscalibration patterns, yielding robust, uncertainty-aware risk estimates suitable for personalized decision support.

Sunday O. Aghamie, Holighi Khalil Shafie
University of Northern Colorado, Greeley, CO, USA
aghamie@gmail.com, khalil.shafie@unco.edu

CP1

Generative Transport for Uncertainty Quantification from Statistical Observations

In many chaotic and high-dimensional dynamical systems, the available data consists only of statistics such as mean, spectra, and higher order moments, providing indirect evidence of the underlying state. Because these statistics do not uniquely determine the underlying dynamics, reconstructing states is inherently uncertain and naturally calls for statistical data assimilation. We address this by introducing a statistic-conditioned generative transport for uncertainty quantification, which models the distribution of plausible short-horizon trajectories while enforcing consistency with observed statistics at every step. At inference, the transport is guided by observed statistics through a gradient signal estimated with Monte Carlo samples of the next-state bridge. This preserves the familiar predict-nudge loop but replaces pointwise likelihoods with statistic likelihoods, producing ensembles that remain consistent with observations and yield calibrated measures of predictive uncertainty without retraining. The method combines statistic-conditioned guidance with higher-order stochastic integration and is evaluated with standard calibration and distance metrics such as coverage, continuous ranked probability score (CRPS), and Wasserstein distance. We demonstrate the approach on Lorenz-63, Lorenz-96, and compressible boundary-layer turbulence, highlighting how different statistic choices affect both parameter/state identifiability and the reliability of uncertainty estimates.

Anubhav Dwivedi
University of Minnesota
anubhavd91@gmail.com

Gs Sidharth
Iowa State University
sidgs@iastate.edu

CP1

Conformalized Ensemble of Neural Networks for Blast Wave Propagation

In this talk, we propose the use of a neural network ensemble to enhance learning and obtain robust confidence intervals on a vector-to-vector regression problem. Relying on the predictive uncertainty decomposition provided by an ensemble of models, we improve training in sparse data regimes through a variance-based active learning strategy. In order to build and calibrate accurate predictive confidence intervals, we apply adaptive conformal prediction based on the ensembles mean and variance. Our approach is evaluated on a challenging nonlinear high-dimensional blast-wave propagation problem, where the goal is to predict the peak overpressure from multi-block input data. Numerical experiments show that a variance-based active learning strategy is particularly well-suited to this setting when the training set is small. By selecting the most informative examples to label, it improves training efficiency and achieves superior performance compared to direct training. Finally, our method produces calibrated intervals that are significantly narrower than those obtained by applying conformal prediction to a single neural network.

Vincent Souveton
CEA, DAM, DIF
vincent.souveton@cea.fr

Pierre Sochala
ASNR (French Authority for Nuclear Safety and Radiation)
pierre.sochala@asnr.fr

Sébastien Terrana
CEA, DAM, DIF
sebastien.terrana@cea.fr

CP1

Latent Generative Modeling of Random Fields from Limited Data

The ability to accurately model random fields plays a critical role in uncertainty quantification (UQ) analyses involving uncertain functional quantities (e.g., heterogeneous material properties, turbulent flows). Deep generative models offer a powerful tool for sampling high- or infinite-dimensional uncertainties like random fields, but their reliance on large, complete training datasets limits their applicability in science and engineering contexts, where data are typically sparse or indirect. In this work, we propose a novel framework for latent generative modeling of random fields that incorporates domain knowledge to supplement limited training data. A variational autoencoder (VAE) is first used to learn compressed, latent representations of functions that adhere to known physical or statistical con-

straints incorporated into the VAE loss function. Generative modeling is then performed in the latent space to yield higher quality random field samples, in particular for complex (e.g., multimodal, heavy-tailed) function distributions. Efficacy is demonstrated on two challenging applications: wind velocity field reconstruction from sparse sensors and material property inference from a limited number of indirect measurements. Results show the effectiveness of incorporating domain knowledge constraints for data-limited problems and the improved performance of the latent generative modeling approach versus directly sampling the pre-trained VAE as in existing work.

James Warner
Nasa Langley Research Center
james.e.warner@nasa.gov

CP2

Generative Physics-Aware Neural Implicit Solvers for Multiscale, Parametric Pdes with Applications in Heterogeneous Media. GenPANIS

Solving parametric Partial Differential Equations (PDEs) is a very challenging task, especially in very high-dimensional cases. In recent years, the advancements in deep learning methods have led the scientific community gradually to more and more complex, sophisticated, and computationally expensive model architectures for capturing the input-output map. These data-driven methods employ architectures that are mostly inspired by successful image-to-image methods. Such models always fail in clearly out-of-distribution scenarios, because they contain no physics-related features in their architecture. We focus on the problem domain of high-dimensional random heterogeneous media, and we propose a lightweight probabilistic generative model that incorporates physics in its architecture in the form of a vectorized, auto-differentiable coarse-grained solver. We demonstrate, clearly, its superior performance against state-of-the-art models with orders of magnitude more parameters in challenging out-of-distribution forward and inverse problems with extremely sparse and noisy data.

Matthaios Chatzopoulos, Phaedon-Stelios Koutsourelakis
Technical University of Munich
matthaios.chatzopoulos@tum.de,
p.s.koutsourelakis@tum.de

CP2

TI-DeepONet: Learnable Time Integration for Stable Long-Term Extrapolation

Reliable long-term forecasting of dynamical systems remains a key challenge for neural operators, which often degrade when extrapolating beyond their training horizon. Standard Deep Operator Networks (DeepONets) either disregard temporal causality through fixed-horizon rollouts or accumulate errors in autoregressive schemes, limiting robustness under uncertainty. We introduce TI-DeepONet, a physics-aware framework that integrates neural operators with adaptive numerical time-stepping. By reformulating the learning task as the approximation of instantaneous time-derivative fields, our method preserves the Markovian structure of dynamical systems, supports continuous-time prediction, and enables higher-order time integrators at inferencebalancing efficiency and accuracy. We further propose TI(L)-DeepONet, which employs learnable coefficients in multi-stage integration schemes to adapt to system-specific variability and enhance predictive fidelity. Across five benchmark PDEs spanning chaotic, dissipative, dispersive, and high-dimensional dynamics, our approach achieves stable extrapolation errors extending to nearly twice the training interval: about 81% lower than autoregressive methods and 70% lower than fixed-horizon rollouts. This research establishes a physics-aware operator learning paradigm that bridges neural approximation with numerical analysis, preserving the causal structure while addressing a critical gap in long-term forecasting of complex physical systems.

Dibyajyoti Nayak, Somdatta Goswami
Johns Hopkins University
dnayak2@jh.edu, somdatta@jhu.edu

CP2

Multi-Fidelity Deep Learning Framework for Parameters Estimations in Hydrodynamics

Despite recent advances in computational power, the computational cost for high-fidelity (HF) simulations remains expensive in many practical problems. Solving inverse problems with HF data alone can be extremely expensive, as many forward simulations must be performed to optimize nonlinear objective functions. In this work, we propose a multi-fidelity (MF) deep learning framework that aims to extract model parameters from experimental observables while lowering the cost of generating HF data and taking advantage of abundant low-fidelity (LF) data. We demonstrate the approach on synthetic inertial confinement fusion (ICF)-like implosion simulations exhibiting Richtmyer-Meshkov instabilities, using a combination of 1D (LF) and 2D (HF) simulations. A surrogate model is first built to create synthetic HF data using many 1D simulations and a few 2D simulations. A variational autoencoder (VAE) is then employed to obtain posterior predictions for the model parameters using synthetic HF data. We demonstrate the ability of the proposed MF framework to accurately estimate the unknown parameters using only a limited amount of 2D simulations in hydrodynamics modeling. The estimations are then validated by plugging in the hydrodynamics code to obtain density fields and observables that satisfy physics consistency. Finally, we demonstrate that this framework achieves both higher accuracy and significant computational savings compared to approaches relying solely on HF data.

Khoa T. Nguyen
Los Alamos National Laboratory
nguyen_k@lanl.gvov

Daniel A. Serino
Los Alamos National Lab
dserino@lanl.gov

Marc Klasky, Balasubramanya Nadiga
Los Alamos National Laboratory
mklasky@lanl.gov, balu@lanl.gov

CP3

Characterizing Collisional Chemistry Uncertainty for Low Temperature Plasma Simulations with Bayesian Inference and Support Vector Machines

Numerically predicting the properties of low-temperature plasma devices, such as torches, relies on characterization of the reaction rates of the constituent plasma species, which depend on the values of electron-impact collision cross sections. These cross sections are epistemically un-

certain in two ways. First, the data obtained for common species transitions through individual experiments and quantum calculations often systematically disagree with each other beyond their measurement errors. Second, most transitions involving higher-lying excited states have no experimental data. To produce reasonable plasma simulation input models to represent these chemistry uncertainties in forward propagation experiments, we use a Bayesian inference model to determine probabilistic models of cross sections for transitions with data, using knowledge of the quantum calculation of electron-argon collisions to enforce appropriate uncertainty bounds through the prior distribution. We then use support vector machines (SVM) to extrapolate these priors to transitions with no data. Using both chemistry-intrinsic properties like the transition threshold energy and extrinsic factors such as the total number of experiments as features, the SVM approach obtains distributions for cross sections of higher-lying state transitions. Given the low amount of data used to train the SVM, we validate these distributions within uncertainty bounds known through domain-specific knowledge.

Garo Bedonian
University of Texas at Austin
Oden Institute
garo.bedonian@gmail.com

Todd Oliver
The University Of Texas-Austin
oliver@oden.utexas.edu

Robert Moser
UT Austin
Oden Institute
rmoser@oden.utexas.edu

CP3

Intelligent Feature Selection in Machine Learning: A Hybrid Optimization Framework

In high-dimensional datasets, feature selection is particularly important for enhancing the computational efficiency, interpretability, and accuracy of machine learning models. The present study delineates the design and development of a feature selection framework that is optimization-based and flexible. This framework incorporates metaheuristic algorithms, including Improved JAYA, PSO, and GA, to identify the most informative feature subsets. The framework is assessed using publicly accessible datasets that encompass behavioural sciences and health informatics. A hybrid objective function that promotes model generalisability is emphasized in our methodology, which balances classification accuracy with feature subset minimization. The proposed approach outperforms traditional filter and wrapper methods in terms of accuracy and reduction ratio when coupled with standard classifiers such as Support Vector Machines (SVM) and Random Forests (RF), as evidenced by experimental results. The proposed framework provides a modular and scalable solution that allows researchers and practitioners to integrate a variety of optimizers based on the nature of their data and problem domain. This work is a valuable contribution to the advancement of machine learning systems that are interpretable, resource-efficient, and high-performing. These systems have a wide range of applications, including medicine diagnosis, stress detection, and human behaviour modelling.

Ramandeep Kumar
Department of Computer Applications,
Lovely Professional University, Phagwara, Punjab, India
raman.pb777@gmail.com

CP3

Continuous-Depth Graph Neural Networks in the Infinite-Node Limit

Continuous-depth graph neural networks, also known as Graph Neural Differential Equations (GNDEs), extend Graph Neural Networks (GNNs) with the continuous-depth architecture of Neural ODEs. What is the behavior of GNDEs as graph size grows? In this talk, we introduce Graphon Neural Differential Equations (Graphon-NDEs) as the infinite-node limit of GNDEs and establish their well-posedness. We present trajectory-wise convergence of GNDEs to Graphon-NDEs as the number of nodes tends to infinity, and we provide explicit convergence rates for GNDEs on both deterministic and random graphs. These results offer theoretical justification for a practical and computationally efficient strategy: deploying GNDE models trained on smaller graphs to much larger ones. Numerical experiments further validate our theoretical findings.

Mingsong Yan
University of California, Santa Barbara
mingsongyan@ucsb.edu

CP4

Modular Jump Gaussian Processes

Gaussian processes (GPs) furnish accurate nonlinear predictions with well-calibrated uncertainty. However, the typical GP setup has a built-in stationarity assumption, making it ill-suited for modeling data from processes with sudden changes, or "jumps" in the output variable. The "jump GP" (JGP) was developed for modeling data from such processes, combining local GPs and latent "level" variables under a joint inferential framework. But joint modeling can be fraught with difficulty. We aim to simplify by suggesting a more modular setup, eschewing joint inference but retaining the main JGP themes: (a) learning optimal neighborhood sizes that locally respect manifolds of discontinuity; and (b) a new cluster-based (latent) feature to capture regions of distinct output levels on both sides of the manifold. We show that each of (a) and (b) separately leads to dramatic improvements when modeling processes with jumps. In tandem (but without requiring joint inference) that benefit is compounded, as illustrated on real and synthetic benchmark examples from the recent literature.

Anna Flowers, Christopher Franck
Virginia Tech
arflowers@vt.edu, chfranck@vt.edu

Mickaël Binois
Inria Center at Université Côte dAzur
mickael.binois@inria.fr

Chiwoo Park
University of Washington
chiwpark@uw.edu

Robert Gramacy
Virginia Tech

rbg@vt.edu

CP4

Bayesian Vecchia Gassian Process Tree Model for High-Dimensional Data

In many real-world applications, Gaussian process (GP) has been widely used for nonparametric regression due to its ability of spatial smoothing via covariance modeling; however, GP suffers from computational challenges and is not well-suited for modeling heterogeneous structures in the data. To address this issue, we develop a Bayesian treed Vecchia GP model by exploiting the ideas in Vecchias approximation and the treed GP model. The proposed model can not only allow efficient computation for large datasets but also allow spatial nonstationarity modeling. To enable fast posterior inference, we also design a sequential Monte Carlo algorithm to generate posterior samples from the tree efficiently, thus avoiding slow-mixing issues of using reversible-jump Markov chain Monte Carlo algorithms. We demonstrate its superior performance in both simulation examples and a real-world application.

Cong Ye
Iowa State University
congye@iastate.edu

Pulong Ma
Department of Statistics
Iowa State University
plma@iastate.edu

CP5

Tensor Decompositions for Interpretable Low-Rank Representation of Atmospheric Data

The Quasi-Biennial Oscillation (QBO) is a quasi periodic wind pattern in the equatorial stratosphere that plays an important role in large-scale variability of the atmosphere. Its structure is high-dimensional, involving dependencies across latitude, longitude, time, and pressure levels. Accurately modeling patterns from such data is challenging due to its scale and complexity. In this work, we investigate tensor decompositions, specifically the Canonical Polyadic Decomposition, as a framework for revealing interpretable low-rank structure in QBO datasets. These methods provide a systematic way to separate modes of variability, quantify contributions across different dimensions, and highlight underlying dynamical processes. We discuss how tensor-based approaches can serve as an effective tool for calibration and uncertainty quantification by enabling reduced-order representations that capture essential variability.

Riti Bahl
Emory University
riti.bahl@emory.edu

Luis Damiano
Sandia National Laboratories
ladamia@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Khachik Sargsyan, Mike Eldred
Sandia National Laboratories
ksargsy@sandia.gov, mseldre@sandia.gov

CP5

A GlobalLocal Multiscale Markov Chain Monte Carlo Method for Domain-Decomposed Bayesian Inference

Uncertainty quantification in PDE-constrained inverse problems is vital to engineering and physical science applications. Yet Markov chain Monte Carlo (MCMC) remains computationally demanding due to high dimensionality and slow multi-chain convergence. The multiscale sampling method (MSM), inspired by domain decomposition for PDEs, accelerates MCMC by exploiting spatial structure. We introduce a globallocal MSM that tackles the curse of dimensionality at two levels: the fine-grid problem is first approximated by a truncated KarhunenLove expansion (KLE); then the global KLE is reduced via local, low-dimensional KLEs over non-overlapping subdomains. This yields faster convergence, typical of lower-dimensional systems. Each iteration combines one global update with parallel local updates, creating a multi-chain scheme well suited to modern HPC. To assess costaccuracy trade-offs, we study the gap between the true posterior and approximate posteriors from the new approach. Results show improved convergence and scalable uncertainty quantification for PDE models.

Aidan R. Gettemy
University of Texas at Dallas
arg210006@utdallas.edu

Felipe Pereira
Department of Mathematical Sciences
University of Texas at Dallas, Dallas, Texas
felipepereira62@gmail.com

Marcio Borges
LNCC/MCTIC
marcio.rentes.borges@gmail.com

CP5

Vector-Valued Surrogate Models Combining Linear Dimension Reduction and Neural Networks

In this talk, we introduce a framework to solve complex multi-block vector-to-vector regression problems. The proposed approach combines supervised or unsupervised linear dimension reduction with nonlinear mapping in latent space using neural networks. A subspace error metric is employed to determine the optimal reduced space dimensionality given the training set size. To learn the mapping in the reduced space, we propose two novel statistically-informed strategies: a variance-informed method that trains a single neural network, and a dependency-informed method that constructs a sequence of neural networks tailored to each output coordinate. Our methods are applied to a blast wave propagation problem which consists in predicting the maximum overpressure with two sources of uncertainty, namely the height of burst and the elevation field both scaled by the explosion yield. Numerical experiments across various training set sizes and reduced space dimensionalities demonstrate the accuracy and robustness of the proposed methods. Compared to high-dimensional mapping with a large neural network, our approach achieves superior performance with significantly reduced computational cost. Furthermore, the resulting surrogate model enables efficient statistical analysis of the predictions, including conditional probability distributions and global sensi-

tivity fields.

Pierre Sochala
ASNR (French Authority for Nuclear Safety and Radiation)
pierre.sochala@asnr.fr

Vincent Souveton, Sébastien Terrana
CEA, DAM, DIF
vincent.souveton@cea.fr, sebastien.terrana@cea.fr

CP5

Surrogate-Aided Bayesian Calibration with Adaptive Learning Strategies

This work investigates an adaptive surrogate model implementation for accelerating Bayesian inference of computationally expensive engineering models. Gaussian Process (GP) regression is adopted as surrogate modelling technique, while the formulation is embedded within a sequential Monte Carlo approach. The adaptive training of the GP is established by iteratively adding training points to inform the target posterior density approximation. At each iteration, a GP model is constructed using the current set of training points to approximate the system response and, consequently, the posterior distribution. This approximation is then used to facilitate sampling from the posterior. The iterative procedure terminates once the posterior density approximations from successive iterations exhibit sufficient similarity. Convergence is assessed using the similarity of the approximated posterior densities. If convergence is not established, the surrogate model is refined through an adaptive DoE that uses the weighted integrated mean squared error as acquisition function. Through proper selection of the weights, exploitation and exploration strategies are promoted to improve, respectively, efficiency and robustness for the GP approximation. Strategies are examined for improving computational efficiency for the DoE optimization and for the sequential Monte Carlo sampling once the surrogate model has been updated.

Alexandros A. Taflanidis
University of Notre Dame
a.taflanidis@nd.edu

Bangalore Satish Aakash
University of California at Berkeley
bsaakash@berkeley.edu

Sang-ri Yi
Rice University
sangri.yi@rice.edu

Joel Conte
University of California at San Diego
jpconte@ucsd.edu

CP6

Movement-Derived Contact Networks and Latent Embeddings for Wildlife Health Modeling and Uncertainty Quantification

Latent embeddings map individual animals, represented as nodes of a graph, to unique vectors containing rich contextual information (e.g., behaviors). These vectors enable us to model the probability of a connection between two nodes as a function of their latent positions in time. In this presentation, we apply a random dot product graph (RDPG) framework to estimate latent positions by exploiting the graph's spectral structure using adjacency spectral embedding. For spatio-temporal data, we extend this idea to dynamic RDPG and temporal graph neural networks to learn time-varying interaction embeddings from evolving networks. In a case study of chronic wasting disease, pathogen transmission occurs through direct and indirect environmental exposure. The challenges from uncertain population sizes, spatial dependence, and changing landscapes complicate the statistical inference of animal movement and contact rates. Using GPS collar data consisting of recorded animal locations at specified intervals in time, we compute the trajectory embeddings and fuse them with interaction embeddings for predicting when, where, and which animals will interact. These embeddings allow us to quantify indirect exposure via visits within a threshold distance to contaminated sites and are critical for downstream tasks such as estimating pairwise contact probabilities and node-level contact rates, while explicitly accommodating spatial structures and environmental dynamics.

Md Mahedi Hasan
University of Wisconsin-Madison
mhasan32@wisc.edu

Daryl DeFord
Dept. of Mathematics and Statistics, Vassar College, NY, USA
ddeford@vassar.edu

Ian McGahan, Juan Francisco Reyes
Dept. of Statistics, University of WisconsinMadison, USA
mcgahani@sou.edu, mandujanorey@wisc.edu

Daniel Storm
Wisconsin Department of Natural Resources, WI, USA
danielj.storm@wisconsin.gov

JUN Zhu
Department of Statistics
University of Wisconsin - Madison
junzhu@wisc.edu

Daniel Walsh
USGS MT Coop Wildlife Research Unit, Univ. of Montana, USA
dan.walsh@mso.umt.edu

CP6

Latent Score-Based Diffusion Model for Data Assimilation in Subsurface Flow

Data assimilation, using multimodal monitoring data, plays a critical role in forecasting and managing large-scale subsurface flow systems. Traditional inversion methods for inferring geological parameters (e.g., permeability) and predicting state variables (e.g., pressure, saturation) often face challenges due to the high-dimensional, nonlinear nature of the problem and the computational cost of ensemble simulations. In this work, we develop a novel data assimilation framework using a latent score-based diffusion model for joint inference of model parameters and state variables. Autoencoders are used to transform both geological models and dynamic state fields into low-dimensional latent space. A score-based diffusion model is then applied in the latent space as a conditional sampler for the joint posterior distribution, enabling efficient assimilation of multi-source observations without retraining the score estimator. We demonstrate this framework on a synthetic 3D heteroge-

neous geological carbon storage model, incorporating both in-situ monitoring well data and interpreted 4D seismic observations. The posterior results show that the method significantly reduces uncertainty in model parameters and predictions of pressure and saturation. We also evaluate the impact of observation type on posterior uncertainty, presenting the potential of this approach for integrating multi-source observations and real-time monitoring planning in subsurface applications.

Su Jiang
Carnegie Mellon University
sujiang@andrew.cmu.edu

CP6

Uncertainty Quantification in Forecasting Effective Motion Models

A central problem in the study of stochastic dynamical systems is forecasting effective motions when the underlying microscopic dynamics are both high-dimensional and computationally expensive. A classic example is molecular dynamics, which generates realistic models of molecular systems but requires simulating the full motion of all atoms, with dimensionality proportional to the number of tracked atoms and often prohibitive cost. Yet, meaningful analysis is usually performed in reduced representations such as principal components or reaction coordinates. We investigate probabilistic time-series forecasting (TSF) methods as a route to obtain reduced dynamical models for such systems. Specifically, probabilistic TSF yields not only point predictions but also full predictive distributions, thereby quantifying uncertainty in the forecast. The resulting distributions approximate effective dynamical models in lower-dimensional coordinates while capturing both deterministic trends and stochasticity. In this talk, I will investigate the application of deep learning tools for probabilistic TSF in molecular dynamics and show how these methods can guide analysis and enable practical applications. By emphasizing predictive distributions, the formulation ensures consistency with the variability structure of the original dynamics and points toward new strategies for efficient analysis of complex stochastic systems.

Pedro Pessoa
Arizona State University
ppessoa@asu.edu

CP6

Flow Matching for Probabilistic Learning of Dynamical Systems from Missing Or Noisy Data

Learning dynamical systems is crucial across many fields, yet applying machine learning techniques remains challenging due to missing variables and noisy data. Classical mathematical models often struggle in these scenarios due to the arose ill-posedness of the physical systems. Stochastic machine learning techniques address this challenge by enabling the modeling of such ill-posed problems. Thus, a single known input to the trained machine learning model may yield multiple plausible outputs, and all of the outputs are correct. In such scenarios, probabilistic forecasting is inherently meaningful. In this study, we introduce a variant of flow matching for probabilistic forecasting which estimates possible future states as a distribution over possible outcomes rather than a single-point prediction. Perturbation of complex dynamical states is not trivial. Community uses typical Gaussian or uniform perturbations to crucial variables to model uncertainty. However, not all variables behave in a Gaussian fashion. So, we also propose a generative machine learning approach to physically and logically perturb the states of complex high-dimensional dynamical systems. Finally, we establish the mathematical foundations of our method and demonstrate its effectiveness on several challenging dynamical systems, including a variant of the high-dimensional WeatherBench dataset, which models the global weather at a 5.625 meridional resolution.

Siddharth Rout, Eldad Haber
University of British Columbia
bnr.siddharthrout@gmail.com, ehaber@eoas.ubc.ca

CP7

Model Error Covariance Estimation for Weak Constraint Variational Data Assimilation

The accuracy of state estimates from weak-constraint four-dimensional variational (4D-Var) data assimilation depend strongly on data and model error covariances. While data error is assumed known, we focus on estimating the model error covariance by framing weak-constraint 4D-Var as a regularized inverse problem, where the inverse model error covariance acts as the regularization matrix. We examine isotropic and non-isotropic forms with hyperparameters including variance, spatial correlation length, and temporal correlation scale. Using the representer method, we reduce the problem from state space to data space, enabling efficient application of regularization parameter selection techniques to estimate covariance hyperparameters. The representer approach also provides analytic expressions for the state estimate, allowing matrix forms for the L-curve, generalized cross-validation (GCV), and Chi-square methods. We validate the approach by assimilating simulated data into a 1D transport equation for wildfire smoke transport under varying noise and model perturbations. Results show the methods successfully estimate hyperparameters for both isotropic and non-isotropic covariances, indicating isotropic forms suffice when the first guess is more accurate than the data, while non-isotropic covariances are preferred when data is more reliable.

Sandra R. Babyale
Boise State University
Boise State University
sandrababyale@u.boisestate.edu

Jodi Mead
Boise State University
Department of Mathematics
jmead@boisestate.edu

Donna Calhoun, Patricia O. Azike
Boise State University
donnacalhoun@boisestate.edu, patriciaazike@u.boisestate.edu

CP7

Error Analysis of the Projected Po Method with Additive Inflation for the Partially Observed Lorenz 96 Model

We consider the filtering problem with the partially observed Lorenz 96 model. Although the accuracy of the 3DVar filter applied to this problem has been established, that of the EnKF has not yet been. This study aims to establish the error bound of a variant of the EnKF, known as the PO method. By introducing the additive inflation and a projection of the background covariance to the ob-

servation space, we establish the error bound of the PO method. A numerical example validates theoretical findings and shows the potential to extend the analysis. We also discuss the relationships between the analyses of the EnKF and those of other data assimilation algorithms.

Kota Takeda
Kyoto University
RIKEN Center for Computational Science
takeda@na.nuap.nagoya-u.ac.jp

CP7

Multiple Data Assimilation As An Approximate Maximum A Posteriori Estimator

This work reinterprets Multiple Data Assimilation (MDA), a recursive update widely used in ensemble Kalman filtering and smoothing, through the lens of maximum a posteriori (MAP) estimation. For linear-Gaussian systems, MAP and minimum mean-square error estimates coincide, but for nonlinear observation models they diverge. Although MDA has been empirically successful, its theoretical connection to MAP has remained unclear. We show that MDA can be derived as a homotopy-based numerical approximation to an ordinary differential equation (ODE) that transports the prior mean to the MAP estimate. This perspective reveals MDA not as an optimization algorithm, but as a first-order ODE solver. Building on this framework, we propose a RungeKutta inspired adaptive-step MDA algorithm with error control. Numerical experiments, including assimilation in the chaotic Lorenz 96 system with nonlinear observations, demonstrate that the new adaptive approach achieves improved accuracy and efficiency compared to standard MDA variants. While MDA is not an exact MAP estimatorsince it omits higher-order termsit nonetheless provides a practical and effective approximation when observation models exhibit sufficient local linearity. This homotopy formulation clarifies MDAs theoretical foundation and motivates extensions toward more accurate and robust ensemble-based filtering methods.

Renato Zanetti
University of Texas at Austin
Department of Aerospace Engineering and Engineering Mechanic
renato@utexas.edu

Andrey A. Popov
University of Hawai'i at Manoa, U.S.
apopov@hawaii.edu

Kristen Michealson
University of Texas at Austin
kmichaelson@utexas.edu

CP8

A Nitsche Approach for Coupled Multiphysics Problems

In this talk, we present a mathematical and numerical model for the interaction between free fluid flow and a poroelastic medium. The Brinkmann equation is used to describe the flow within the porous domain, allowing for the inclusion of inertial effects. A thermodynamically consistent framework is proposed for modeling soft tissue perfusion. To approximate the solution, we develop a mixed finite element method based on Nitsches approach and establish the well-posedness of the discrete formulation. Furthermore, we derive a priori error estimates in the energy norm. Numerical experiments are provided to confirm the theoretical convergence rates and to demonstrate the methods capability in accurately capturing the underlying physical behavior.

Aparna Bansal
Indian Institute of Technology Roorkee
a_bansal@ma.iitr.ac.in

Nicolas A. Barnafi
Pontificia Universidad Católica de Chile
nicolas.barnafi@uc.cl

Dwijendra Narain Pandey
Indian Institute of Technology Roorkee
dwijpfma@iitr.ac.in

Ricardo Ruiz-Baier
Monash University
ricardo.ruizbaier@monash.edu

CP8

Preconditioned CrankNicolson Differential Evolution Markov Chain Monte Carlo Method Applied to Stochastic Two-Phase Flow Problem in Porous Media

The shape and size of the proposal distribution are known to be very crucial for the convergence of the Markov chains, mainly in porous media fluid flows, where the stochastic dimension can be high. In this work, we use the new Preconditioned CrankNicolson Differential Evolution Markov Chain Method (pCN-DE) in a two-phase flow problem. The pCN-DE method is a population-based approach that combines the classical Preconditioned CrankNicolson (pCN) and Differential Evolution (DE) Metropolis algorithms, allowing for the exchange of information between the parallel chains and providing information about the scale and orientation of the jumping distribution with fewer limitations with respect to the number of parallel chains needed. On the other hand, the pCN contribution guarantees the invariance of the proposals. The new method yielded significantly better results compared to the classical pCN.

Marcio Borges
LNCC/MCTIC
marcio.rentes.borges@gmail.com

Felipe Pereira
Department of Mathematical Sciences
University of Texas at Dallas, Dallas, Texas
felipepereira62@gmail.com

Michel Tosin
LNCC/MCTI
michelt@posgrad.lncc.br

CP8

Implicit Time Integration of High-Dimensional Tensor Differential Equations on Tensor Train Manifolds Based on a Cross Algorithm

This work presents an algorithm for the implicit time integration of high-dimensional tensor differential equations (TDEs) on tensor train manifolds. In addition to time-dependent problems, the framework also applies to linear tensor equations. The method offers several advantages:

(i) it achieves substantial savings in memory and floating-point operations by employing a cross-approximation strategy based on CUR decomposition combined with the discrete empirical interpolation method, (ii) it handles TDEs with general nonlinear terms, regardless of their specific form, and (iii) it incorporates rank adaptivity, enabling the tensor rank to adjust automatically to meet a desired accuracy. The algorithms performance is demonstrated through a range of numerical examples.

Behzad Ghahremani, Hessam Babaee
University of Pittsburgh
beg66@pitt.edu, h.babaee@pitt.edu

CP8

A High-Dimensional Density Estimation Method and Its Application for Solving PDEs

Probability density estimation remains an open challenging problem in computational science and engineering. By coupling the Knothe-Rosenblatt (KR) rearrangement and the flow-based generative model, we developed an invertible transport map, called KRnet, for high-dimensional density estimation. In this talk, we give an overview of KRnet and discuss its adaptive version for the Fokker-Planck equations and stochastic dynamic systems. KRnet for solving Bayesian inverse problems is also studied. This is joint work with Yani Feng, Junjie He, Kejun Tang and Xiaoliang Wan.

Qifeng Liao
ShanghaiTech University
liaoqf@shanghaitech.edu.cn

CP9

Bayesian Calibration with Functional Outputs Using Elastic Partial Matching

Calibrating a simulation model involves estimating model parameters by comparing its outputs with experiences to ensure that simulation results accurately reflect these experiences. However, when outputs are functions of time, there are multiple ways to quantify the difference between experimental and simulated outputs. It has recently been proposed to use elastic functional data analysis, which makes it possible to decompose a functional output into two new functions: a function temporally aligned to a template and the corresponding warping function. This approach splits the problem into two independent calibration tasks to address functional misalignment. The first task is based on the comparison of the aligned functions and the second one is based on the comparison of the warping functions. In this work, we extend this framework with elastic partial alignment and with the use of Gaussian process regression. We apply the method to a test case and compare it to a Bayesian calibration method that does not use functional alignment.

Paul Casteras
CEA, DAM, DIF, F-91297 Arpajon, France
paul.casteras31@gmail.com

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Julien Bect
Université Paris-Saclay, CNRS, CentraleSupélec
Laboratoire des signaux et systèmes
julien.bect@centralesupelec.fr

Gwenael Salin
CEA, DAM, DIF, F-91297 Arpajon, France
gwenael.salin@cea.fr

CP9

Exact Uncertainty Propagation Through Solvers for Linear Elliptic Partial Differential Equations

We present a novel framework for deriving push-forward measures for stationary, multigrid, and domain-decomposed solvers for linear elliptic partial differential equations of the form

$$\nabla_x \cdot (a(x)\nabla_x u(x)) = f(x,\omega),$$

where f is a random field. By exploiting the linearity of these solvers and the closure of Gaussian random variables under linear transforms, our method efficiently propagates uncertainties. We relax Gaussianity assumptions using the denseness of Gaussian mixture models in the space of continuous probability measures, allowing arbitrarily accurate approximations of complex distributions via variational inference. Benchmarks validate our approach, comparing uncertainty propagation in diverse linear systems against Monte Carlo simulations. Our framework extends beyond elliptic boundary value problems, enabling robust and precise uncertainty quantification for a wide range of linear algebraic systems. These results underscore its potential for impactful applications in computational science and engineering.

Rahul Manavalan
Chair for Scientific Computing in Computer Science
Technical University of Munich
rahul.manavalan@matstat.lu.se

CP9

Information Geometry of Uncertainty Quantification for Downstream Decision Making: Resource Theory Perspective

Consider Decision Maker (DM) attempting to maximize utility $u(x,y)$ over decision x, where random variable y represents uncertain State of Environment (SoE). Bayesian DM maximizes the expected utility assuming known probability distribution $Q_(y)$. Robust Bayesian DM assumes that distribution of y, Q(y) lies within certain Ambiguity Set (AS) P and maximizes the expected utility against the worst-case scenario distribution Q within this set. If DM is capable of decision randomization, this optimization defines Phi-equilibria in zero-sum game between DM and Adversarial Environment (AE) where P represents the AE's resources. We propose Resource Theory based Uncertainty Quantification with respect to the ability of the AE's resources to cause loss in the expected DM utility by manipulating AS P. Given prior $Q_(y)$, we consider three types of AE's uncertainty: Agnostic, Distributional, and Entanglement, with ASs P(H), P(H,D), and P(H,D,I) respectively. P(H) is comprised of distributions $Q(y)\ exp[(lnQ_(y))/T], T > 0$ with entropy not exceeding H. P(H,D), allowing AE to benefit from knowledge of the DM utility, is comprised of distributions Q with entropy not exceeding H and Kullback-Leibler deviation from P(H) not exceeding D. Finally, P(H,D,I), allowing AE's strategies correlated with DMs decisions, is comprised of conditional distributions $Q(y|x)$ for which mutual information

between y and x does not exceed I and joint distribution Q(x,y) satisfies the above constraints.

Vladimir Marbukh
National Institute of Standards and Technology
marbukh@nist.gov

CP10

A Cur Krylov Solver for Large-Scale Generalized Sylvester and Non-Sylvester Linear Matrix Equations

This paper introduces a novel methodology leveraging CUR decomposition for solving large-scale generalized Sylvester as well as non-Sylvester equations on low-rank matrix manifolds. The approach decomposes the original equations into two smaller subproblems, each characterized by substantially reduced coefficient matrices. This decomposition facilitates the application of existing efficient solvers specifically designed for the structures where the unknown matrix is large from one side and small from the other side. By selectively solving the equation at strategically chosen rows and columns instead of the full matrix, the proposed method significantly reduces computational overhead. Additionally, the method incorporates rank adaptivity, dynamically adjusting computational rank to reach the desired accuracy. A Krylov-based scheme is presented to solve the resulting subproblems, which enables solving large-scale problems and does not require Sylvester structure. The methodology is demonstrated in three settings: (i) implicit time integration of matrix differential equations on low-rank manifolds, (ii) steady-state Lyapunov equations, and (iii) non-Sylvester linear matrix equations, such as those arising in nonlinear matrix problems.

Saeed Akbari, Hessam Babaee
University of Pittsburgh
s.akbari@pitt.edu, h.babaee@pitt.edu

Damiano Lombardi
INRIA Paris-Rocquencourt
damiano.lombardi@inria.fr

CP10

Likelihood-Tempering Sequential Monte Carlo Method and Its Applications to Chemical Engineering Problems

In this study, we propose and demonstrate a sequential Monte Carlo (SMC) method incorporating likelihood tempering for nontime-series data, which accelerates uncertainty quantification when the model is expensive to evaluate. Unlike Markov Chain Monte Carlo (MCMC) methods that rely on sequential steps, or conventional SMC methods where time-series measurements are assimilated sequentially, the proposed likelihood-tempered SMC modifies the likelihood function sequentially and approximates it using scattered particles. The algorithm can be parallelized, as the computation of particles can be distributed across multiple processors. The advantages of the likelihood-tempered SMC are demonstrated with multiple examples. Its performance over conventional MCMC methods is quantified in several benchmark problems. In addition, the approach is applied to a real chemical engineering problem, where parameters in a complex differential equation model are estimated and their uncertainties quantified. Furthermore, computational time is significantly reduced through parallel implementation.

Tatsuki Maruchi, Tomoyuki Yajima, Yoshiaki Kawajiri
Nagoya University
Materials Process Engineering
maruchi.tatsuki.t7@s.mail.nagoya-u.ac.jp,
yajima.tomoyuki@material.nagoya-u.ac.jp,
kawajiri@nagoya-u.jp

CP10

Efficient and Scalable Method for Sparse Rayleigh Quotient Maximization

Sparse Rayleigh Quotient Maximization (sparse RQM) is a cardinality-constrained optimization problem that arises in many contexts, including machine learning, statistical analysis, dimensionality reduction, and dynamical systems analysis. The cardinality constraint in sparse RQM turns the standard RQM into an NP-hard problem, requiring a combinatorial search to find globally optimal solutions. Here, we introduce an efficient iterative method that - in contrast with L1-regularization - provides both a priori guarantees on the degree of solution sparsity and performance bounds on the optimization objective, computing (sub)optimal solutions to the sparse RQM problem. We use several benchmark problems to demonstrate the method's cubic convergence and scalability to problem-instances with thousands of decision variables. We then apply our method in the context of dynamical systems analysis and reduced-order modeling of a fluid flow, wherein our approach successfully identifies sparse and spatially localized instability mechanisms that drive energy amplification in a compressible turbulent boundary layer.

Jonathan Smith, Talha Mushtaq
University of Minnesota
smi01444@umn.edu, talhamushtaq15@gmail.com

Maziar S. Hemati
Aerospace Engineering and Mechanics
University of Minnesota
mhemati@umn.edu

CP10

Nomad for Stochastic Simulator and One Shot Uncertainty Quantification

We propose a neural operator framework, termed mixture density nonlinear manifold decoder (MD-NOMAD), for emulating the conditional response distribution of stochastic simulators. Our framework leverages an amalgamation of the pointwise operator learning neural architecture, nonlinear manifold decoder (NOMAD), with versatile mixture density-based methods for approximating the conditional probability distributions of stochastic output functions. Furthermore, MD-NOMAD harnesses the ability of probabilistic mixture models to estimate complex probability distributions and inherits the high-dimensional scalability of NOMAD. A key feature of MD-NOMAD is its ability to perform one-shot uncertainty propagation, eliminating the need for Monte Carlo methods in statistical computations for uncertainty propagation tasks. We conduct empirical assessments on a wide array of stochastic ordinary and partial differential equations and present the corresponding results, which highlight the performance of the proposed framework.

Akshay Thakur
University of Notre Dame

athakur3@nd.edu

Souvik Chakraborty
Department of Applied Mechanics
Indian Institute of Technology Delhi
souvik@am.iitd.ac.in

CP11

Comparative Analysis of Algorithmic Performance and Bias Metrics for Bias Detection: A Pre and Post-Fairness Intervention Study

Fairness in machine learning systems is a major concern in sensitive AI domains. Biased outcomes generally emanate from imbalanced data that affect protected groups. This study assesses two approaches in detecting bias, algorithm performance and fairness specific bias metrics to determine the most effective approach in detecting and quantifying bias, and to assess the bias mitigation strategy. The proposed bias mitigation strategy targets the root causes of data imbalance that mainly contribute to such disparate outcomes. Quantifying bias involves the use of detection metrics like Disparate Impact, Demographic Parity, and Equal Opportunity Difference, and also involving algorithmic performance metrics such as accuracy, precision, recall, and F1 score. This allows for a comparison across different algorithmic architectures with respect to fairness. Results showed that 1) the fairness intervention reduced bias, as shown by a significant decrease in bias metrics, 2) the gained fairness had no cost on algorithm performance, and 3) the comparative analysis showed that bias metrics are more sensitive to disparities detection, while algorithmic performance provides insights into the overall reliability of the models. We may conclude that this work stresses the importance of including both fairness-specific and traditional performance metrics in bias detection toward a comprehensive assessment of the effect of fairness interventions

Muna Abdelrahim
Jackson State University
muna.abdelrahim@students.jsums.edu

CP11

Rnyi Extropy for Uncertainty Quantification: Theory and Applications

This paper revisits the framework of Rnyi extropy and extends its mathematical properties for uncertainty quantification. New results include its convergence to information extropy at limiting values and its ability to assume both positive and negative measures, providing deeper insight into uncertainty. We also unify Rnyi extropy with information and Tsallis extropies. To illustrate practical relevance, we apply Rnyi extropy to cryptocurrency markets, showing its advantage in capturing non-Gaussian dynamics and quantifying risk more effectively than traditional entropy-based measures. Further, we integrate Rnyi extropy with machine learning models for price prediction, highlighting its role as a versatile tool for uncertainty quantification in complex systems.

Luckshay Batra
BML Munjal University
Kapriwas, Haryana 122413
luckshay25@gmail.com

Ruchika Lochab
International Management Institute
Qutab Institutional Area, New Delhi, Delhi 110016, India
rchklchb@gmail.com

CP11

Ellipsoidal State Estimation for Quadratic Systems with Unknown But Bounded Noise

This talk addresses state estimation for quadratic dynamical systems with linear measurements in the presence of unknown but bounded process disturbances and sensor noise. We use Lyapunov theory and Petersens lemma from the robust controls literature to establish convex semi-definite programming conditions that yield an estimator with guaranteed ellipsoidal error bounds. In contrast to Kalman and particle filtering approaches, the ellipsoidal viewpoint delivers worst-case, certificate-backed assurances without relying on noise distributions, making it suitable when distributions are unavailable or unreliable. We demonstrate our approach on a series of examples, including inertial navigation, aerodynamic flow reconstruction, and joint state-parameter estimation.

Amir Enayati Kafshgarkolaei
University of Minnesotaota
University of Minnesotaota
Enaya004@umn.edu

Maziar S. Hemati
Aerospace Engineering and Mechanics
University of Minnesota
mhemati@umn.edu

CP12

Pac-Bayes Certificates for Bayesian Inverse Problems: A Case Study on the Heat Equation

Although Bayesian inverse problems for partial differential equations (PDEs) yield posterior distributions and credible intervals, these can be overconfident in sparse or noisy sensing regimes. We use the one-dimensional heat equation as a test case to introduce PAC-Bayes certificates into the Bayesian inversion framework. Our approach connects tempered (Gibbs) posteriors with PAC-Bayes theory to obtain finite-sample, distribution-free guarantees, extended with a mesh-robust decomposition that separates discretization from generalization. The reliability of PAC-Bayes certificates in detecting overfitting is shown by experiments with varying sensor counts, noise levels, and discretizations. Gaps between the certified bound B_λ and the empirical risk $\widehat{L}$ identify high-complexity cases ($m = 5$, small n) and weak regularization (e.g., $\lambda = 0.5$). Interpretability is improved by the discretization penalty η_h, which isolates solver error. A grid of 1,728 experiments was completed in four days on a standard laptop, with each run taking 3-5 minutes, showing computational feasibility. Adaptive MCMC halved sampling while maintaining effective sample size, and certificate computation added under 5% overhead. The method scales with problem size and sensor count, extending beyond the 1D case. PAC-Bayes certification advances reliable, reproducible uncertainty quantification while adding a conservative, auditable layer at negligible cost.

Tanisha Gupta, Ritu Arora
University of Delhi
Janki Devi Memorial College

tanisha.gupta@allpay.net, ritu@jdm.du.ac.in

CP12

Practical Aspects of the Virtual Noise Convex Optimum Design Approach for Correlated Responses

We present several practically-oriented extensions and considerations for the virtual noise method in optimal design under correlation. First we introduce a slightly modified virtual noise representation which further illuminates the parallels to the classical design approach for uncorrelated observations. We suggest more efficient algorithms to obtain the design measures. Furthermore, we show that various convex relaxation methods used for sensor selection are special cases of our approach and can be solved within our framework. Finally, we provide practical guidelines on how to generally approach a design problem with correlated observations and demonstrate how to utilize the virtual noise method in this context in a meaningful way.

Markus Hainy
Johannes Kepler University Linz
markus.hainy@jku.at

Werner Müller
Johannes Kepler University
Austria
werner.mueller@jku.at

Andrej Pázman
Comenius University Bratislava
andrej.pazman@fmph.uniba.sk

CP12

Maximum Entropy Sequential Design for Bayesian Hierarchical Gaussian Processes: Application to Building Physicochemical Surrogate Models

We present an information-theoretic sampling framework for surrogate modeling in hybrid settings, where known physics combines with data-driven inference of a latent function. We place a Gaussian process (GP) prior over the latent function and place priors on the GP hyperparameters to encode structural information from practitioner knowledge. To make inference and acquisition tractable, we approximate the GP predictive posterior with a finite, uniformly weighted mixture of GPs. Our main contribution is the derivation of entropy-based acquisition functions for the mixture representation. We provide a closed-form expression for the differential entropy of the approximate posterior and establish a lower bound using a Taylor expansion around the mean function of each GP component. This contribution yields a principled criterion for selecting input locations that maximize expected information gain. Numerical experiments show that entropy-guided sampling improves sample efficiency by concentrating queries in regions of high predictive uncertainty and information content. The resulting surrogates converge more quickly, achieve higher predictive accuracy, and remain consistent with the structural priors encoded in the GP hierarchy. More broadly, the framework demonstrates how entropy approximations for Gaussian mixture posteriors enable uncertainty-aware acquisition in Bayesian nonparametric modeling.

Kyla D. Jones, Alexander Dowling
University of Notre Dame
kjones29@nd.edu, adowling@nd.edu

CP13

Intrusive Uncertainty Quantification and Optimal Experiment Design to Characterize Membrane Separation Systems

Membrane separation systems are crucial to sustainable water treatment, chemical processing, and the recovery of critical minerals. However, their design is hindered by uncertainty in transport models and parameters. This study introduces an intrusive uncertainty quantification (UQ) and model-based design of experiments (MBDoE) framework to characterize neutral and charged nanofiltration membranes. By leveraging advances in dynamic diafiltration apparatuses, model-based data analytics, and Pyomo.DOE, this framework simultaneously discriminates between transport models, parameter estimates, and selects operating conditions that maximize information content while minimizing experimental effort. Intrusive approaches leverage Fisher information matrix criteria (A-, D-, E-, and ME-optimality) and weighted regression for heterogeneous data, enabling robust parameter estimation and model discrimination in multi-salt, charged transport systems. Finally, the study explores opportunities to create self-driving laboratories (SDLs) and digital twins that autonomously design, execute, and refine experiments. This work envisions establishing a scalable, predictive foundation for uncertainty-aware membrane characterization and inverse design of separation processes in the domains of water, energy, and critical materials.

Keshav Kasturi Rangan, Molly Dougher, Laurianne Estrada, Shammah Lilonfe, Pratham Singh
University of Notre Dame
kkasturi@nd.edu, mdoughe6@nd.edu, llair@nd.edu, slilonfe@nd.edu, psingh4@nd.edu

William Phillip
Department of Chemical and Biomolecular Engineering
University of Notre Dame
wphillip@nd.edu

Alexander Dowling
University of Notre Dame
adowling@nd.edu

CP13

Intrusive Uncertainty Quantification and Optimal Experiment Design in the Open-Source Pyomo Ecosystem

This contribution describes ParmEst [1] and Pyomo.DoE [2], two pillars of the open-source Python-based Pyomo [3] ecosystem for computational optimization with (partial differential) algebraic equation mathematical models. Specifically, ParmEst facilitates intrusive frequentist parameter estimation (PE) and uncertainty quantification (UQ) through built-in features, such as covariance matrix estimation, bootstrapping, and likelihood ratio tests. Complementarily, Pyomo.DoE enables optimal experiment design by maximizing various metrics of the Fisher information matrix, such as A-optimality (trace), D-optimality (determinant), E-optimality (minimum eigenvalue), and ME-optimality (condition number). ParmEst and Pyomo.DoE can solve high-dimensional optimization problems by leveraging the model structure and exact derivative information. Finally, we will discuss future opportunities to integrate PE and UQ capabilities with optimization under

uncertainty, including robust optimization (RO) with nonconvex models via PyROS [4], a Pyomo-based RO package. References: 1. Klise, K. A.; Nicholson, B. L.; Staid, A.; Woodruff, D. L. Comput. Aided Chem. Eng. 2019, 47, 4146. Elsevier. 2. Wang, J.; Dowling, A. W. AIChE J. 2022, 68 (12), e17813. 3. Hart, W. E.; Watson, J.-P.; Woodruff, D. L. Math. Program. Comput. 2011, 3 (3), 219260. 4. Isenberg, N. M.; Akula, P.; Eslick, J. C.; Bhattacharyya, D.; Miller, D. C.; Gounaris, C. E. AIChE J. 2021, 67 (5), e17175.

Shammah Lilonfe, Shuvashish Mondal
University of Notre Dame
slilonfe@nd.edu, smondal@nd.edu

Stephen Cini
University of Notre Dame
Chemical and Biomolecular Engineering
scini@nd.edu

Daniel J. Laky
University of Notre Dame, U.S.
dlaky@nd.edu

Alexander Dowling
University of Notre Dame
adowling@nd.edu

CP13

Bayesian Experimental Design for History Matching

History matching (HM) is a type of computer model calibration that rules out input parameter settings that are inconsistent between observations and computer model output according to a distance-based implausibility function and uncertainty specifications. The remaining parameter space is termed as Not Ruled Out Yet (NROY). History matching has proven to be more effective when performed in waves. At each wave of HM, a new ensemble is obtained to update an emulator (surrogate model) before finding an NROY space. In this talk, we propose a Bayesian experimental design to optimise the data collection process using a loss function that compares the volume of the NROY space obtained with an updated emulator to the volume of the true NROY space obtained using a perfect emulator. The intuition behind the proposed loss function corresponds to the aim of history matching. In particular, by performing history matching, we want to ensure that we are not going to incorrectly rule out points that are in fact close to observations, or leave regions of space that give output far from the observations. We compare the proposed methodology to space-filling design approaches using toy models as well as applications in climate and healthcare models.

Victoria Volodina
Department of Mathematics and Statistics
University of Exeter
v.volodina@exeter.ac.uk

CP13

Bayesian Optimal Experimental Prediction Design for Physical Experiments Informed by Computer Experiments

We consider Bayesian optimal experimental design for prediction at specified design configurations of physical systems. A computer model describing the system is assumed to be available, providing predictions of the output of interest as a function of the design variables and calibration parameters. Since the computer model does not perfectly represent the system, model discrepancy must be taken into account. Moreover, the calibration parameters that best align the model with reality are unknown. To address these uncertainties, we adopt the framework of Kennedy and OHagan (2001), which models both computer experiments and model discrepancy using Gaussian processes. As a design criterion, we use the mutual information between observations at design points and predictions at target points, thereby incorporating parameter uncertainty. We further introduce a new criterion that combines mutual information with a complexity penalty, guiding the choice of design points and reducing oversampling in regions with high predictive uncertainty. We compare this approach with other criteria such as integrated mean squared prediction error and designs that ignore parameter uncertainty. Our examples include a numerical case study and a biochemical dataset from Swameye et al. (2003), considering both sequential offline and adaptive online designs. Finally, we plan to leverage modern machine learning methods to develop non-myopic adaptive designs for this framework.

Hao Zhu
Johannes Kepler University Linz
hao.zhu@jku.at

CP14

A Hierarchical Bayesian Approach for Quantification of Fabrication Uncertainties in Integrated Photonics

Photonic integrated circuits enable high-bandwidth signal processing through carefully engineered interference of optical waves. However, as a consequence, these devices are exceptionally sensitive to even small stochastic variations in the fabrication process; nanometer-scale deviations in width, height, sidewall angle, and roughness of waveguides perturb effective indices and coupling coefficients, resulting in phase errors and spectral shifts that accumulate throughout the circuit. We will discuss the use of hierarchical Bayesian modeling and inference as a principled way to quantify the distributions of process-level fabrication uncertainties in integrated photonics. Accurate electromagnetic forward models are used to map fabrication parameters to device behavior; by inverting the measurements of an ensemble of fabricated devices, we then disentangle approximations of process-level uncertainty distributions from measurement noise and model errors. We contend that such an explicit quantification of fabrication variations is a key step to herald the next generation of robust photonic fabrication; it feeds back naturally into the engineering iteration cycle by informing fabrication process modifications that reduce variations or by suggesting device designs that increase resilience to irreducible fabrication variations.

Flemming Holtorf, Frank Schäfer, Jacob Taylor, Dirk Englund
Axiomatic_AI
holtorf@mit.edu, franks@axiomatic-ai.com, jake@axiomatic-ai.com, dirk@axiomatic-ai.com

CP14

Surrogate-Based Probabilistic Fatigue Assessment of Transmission Lines with GPR: From Damage Es-

timation to Inspection Planning

Transmission line towers are repeatedly subjected to fluctuating wind loads that drive fatigue damage in bolted connections, threatening long-term reliability. This study develops a surrogate-based probabilistic framework to assess fatigue damage and guide inspection planning. High-fidelity finite element simulations are performed to capture joint-level responses under both preloaded and loosened bolt conditions, with fatigue quantified using rain-flow cycle counting and Miner's rule. The resulting dataset is used to train Gaussian Process Regression (GPR) surrogates, which provide rapid and accurate predictions of fatigue accumulation across diverse wind speeds and directions. By coupling these surrogates with joint probability distributions of typical winds and a Poisson-based extreme-event model, the framework synthesizes annual damage and propagates uncertainty through sampling-based simulations. Results highlight fatigue-critical connections, reveal the relative roles of typical versus extreme winds, and quantify the influence of stress concentration factors and bolt preloading states on vulnerability. Building on these insights, inspection strategies are optimized by incorporating detection probabilities, repair resets, and cost–risk trade-offs. The proposed methodology delivers a computationally efficient, reliability-informed tool for fatigue risk assessment and inspection scheduling of transmission line systems.

Pooria Mazaheri
Ph.D. student
Iowa state university
mazaheri@iastate.edu

CP14

Bayesian Optimization of Plasma Kinetic Reaction Conditions

Low-temperature plasma processes are increasingly used in chemical synthesis, energy conversion, and environmental applications. Yet, their systematic study remains difficult because of nonlinear pathways, high-dimensional parameter spaces, and the prohibitive cost of kinetic simulations. To address these challenges, we present a data-driven framework that integrates Gaussian Process (GP) regression with Bayesian Optimization (BO) to efficiently explore and optimize plasma kinetic systems. Initial data from ZDPlasKin, a widely used high-fidelity plasma kinetics simulator, is employed to train GP surrogates that capture complex inputoutput relationships while providing quantified predictive uncertainty. We compare multiple GP modeling choices, including kernel families and hyperparameter learning strategies, and evaluate their effect on predictive accuracy and uncertainty calibration. Building on these surrogates, BO is used to adaptively select new simulation points, and we examine the impact of different acquisition functions (e.g., Expected Improvement, Upper Confidence Bound, Probability of Improvement) on optimization efficiency. This iterative, UQ-driven design reduces the number of simulations needed to map the design space, accelerates identification of high-yield regimes compared to random or space-filling approaches, and demonstrates how uncertainty-aware surrogate modeling enables more systematic and efficient exploration of complex reacting systems.

Vivek Saini, Alexander Dowling, Denver Haycock, William Schneider
University of Notre Dame
vsaini@nd.edu, adowling@nd.edu, dhaycock@nd.edu, wschnei1@nd.edu

CP14

Mixed-Order Adsorption Kinetics under Uncertainty

Adsorption is a fundamental process with wide-ranging applications in catalysis, separations, and environmental protection. While pseudo-first-order and pseudo-second-order models are commonly used to describe adsorption kinetics, complex systems often involve competing mechanisms. The mixed-order model provides a more flexible and accurate description by combining both driving forces. Traditional parameter estimation methods for this model frequently yield biased or unreliable results. To address this, we propose a probabilistic framework that explicitly accounts for uncertainty arising both from the underlying dynamical system and from experimental measurement errors. In particular, we treat the model parameters as continuous random variables and apply the Random Variable Transformation technique to derive the first probability density function of the solution. This approach not only yields distributions for the model parameters but also provides a complete probabilistic characterization of the solution itself, offering a more robust and informative representation of adsorption dynamics. We also highlight ongoing work exploring alternative formulations, such as embedding the deterministic model into a Gaussian Process framework combined with particle filtering.

Sorina M. Sferle
Universitat Politècnica de València
smsferle@upv.edu.es

Juan Carlos Cortés
Universitat Politècnica de València
Instituto Multidisciplinar de Matemáticas
jccortes@mat.upv.es

Oksana A. Chkrebtii
The Ohio State University
oksana@stat.osu.edu

Ana Navarro Quiles
Universitat Politècnica de València
Instituto Multidisciplinar de Matemáticas
annaqui@upv.es

CP15

Performance of Ellipsoidal Estimation for Low-Earth Orbit Determination under Bounded Parametric Uncertainty

This work focuses on an important problem in relation to space situational awareness: determining orbits of objects in low-Earth orbit (LEO) with imprecise knowledge of parameters. State-of-the-art nonlinear estimators like the extended and unscented Kalman filters (EKF and UKF, respectively) have traditionally been utilized for orbit determination; however, accommodating parametric uncertainty poses a non-trivial challenge. To this end, we employ an ellipsoidal state estimation framework (Bhattacharjee and Subbarao, IEEE Trans. Automat. Contr. 67 (2021), pp. 894-901) that considers ellipsoidal characterizations of the process, measurement, and initial condition uncertainties. The framework relies upon pseudo-linear descriptions of the nonlinearity in dynamical and observation models, and it inherits the recursive predictor-corrector structure

of the standard Kalman filter variants. At every prediction and correction step, optimal ellipsoids are computed by solving convex semi-definite programming optimization problems that can be reliably solved using standard off-the-shelf solvers. Numerical experiments are performed to quantify nominal robustness of the EKF and UKF against parameter mismatch for orbit determination of a LEO satellite. The ellipsoidal framework exhibits a higher degree of nominal robustness compared to EKF and UKF. The joint parameter and state estimation problem is also studied for the same orbit, and performance of the three methods are compared.

Diganta Bhattacharjee
University of Minnesota, Twin Cities
dbhattac@umn.edu

Maziar S. Hemati
Aerospace Engineering and Mechanics
University of Minnesota
mhemati@umn.edu

CP15

Stochastic Optimal Control for Trajectory Optimization in Mine Counter Measure Simulations

Simulating mine counter measure (MCM) search missions at sea performed by autonomous vehicles presents a complex challenge. The goal of MCM simulations consists of computing trajectories for autonomous vehicles surveying a domain where sea mines are suspected. These trajectories have to be such that the residual risk of not detected the sea mines after survey is below a certain threshold. In [1], we model the MCM problem as a stochastic optimal control problem. Mathematically, this translates as minimizing the total mission time needed to survey a domain Ω for a given residual risk, i.e.,

$$\min T_F, \tag{1}$$

subjected to

$$\mathbb{E}[q\,(T_F)] := \int_{\Omega} \mathrm{e}^{-\int_0^{T_F} \gamma(x(\tau),\omega)\,d\,\tau} f_X\,(\omega)\,d\,\omega \leq \text{Residual Risk}. \tag{2}$$

In [1], we present a relaxation strategy combined with a quasi-Monte Carlo scheme for solving the integral, which shows a speedup of a factor 2 with respect to Monte Carlo. We also expand our implementation to accommodate generalized non-convex domains, e.g., a hexagon. [1] Blondeel, P., Van Utterbeeck, F., Lauwens, B. (2025). Application of quasi-Monte Carlo in Mine Countermeasure Simulations with a Stochastic Optimal Control Framework, in 16th International Conference on Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing, Waterloo, Canada, 2024.

Philippe Blondeel, Filip Van Utterbeeck, Ben Lauwens
Belgian Royal Military Academy
philippe.blondeel@mil.be, filip.vanutterbeeck@mil.be, ben.lauwens@mil.be

CP15

Uncertainty Quantification in Soft Biological Tissue Growth and Remodeling

Cardiovascular biomechanics often involves variations in tissue mass (growth) and microstructure (remodeling), driven by mechanical or chemical stimuli. The classical constrained mixture theory describes this growth and remodeling (G&R) process by modeling the synthesis and degradation of key extracellular matrix constituents: elastin, smooth muscle cells, and collagen fibers. Accurately characterizing this complex mechanobiological phenomenon is essential for the prevention and diagnosis of vascular diseases, yet it remains challenging due to the large number of model parameters, high computational demands, and potential discrepancies with experimental observations. In this work, we investigate uncertainties in the mechanobiologically equilibrated and time-dependent G&R process, focusing on solution behavior classification and parameter sensitivity analysis. In particular, we propose a framework to efficiently classify the outcomes of G&R simulations under perturbations to material integrity and mechanosensitive parameters. This includes leveraging a flexible neural network for fast evaluation of classification probabilities and extraction of decision boundaries, combined with a variance-based global sensitivity analysis to identify the key parameter components driving transitions in simulation outcomes. Our study utilizes both reduced-order and three-dimensional finite element G&R models, offering novel insights into the understanding of mechanobiological equilibrium, numerical stability and the progression of thoracic aortic aneurysms (TAA).

Guoxiang Tong
Stanford University
gxus34@stanford.edu

Erica Schwarz, Jay Humphrey
Yale University
erica.schwarz@yale.edu, jay.humphrey@yale.edu

Alison Marsden
School of Medicine
Stanford University, USA
amarsden@stanford.edu

CP15

Advancing Storm Surge Hazard Emulation Through the Implementation of Graph Neural Networks

Probabilistic storm surge hazard estimation plays a key role in coastal resilience planning. Uncertainty quantification is achieved by generating an ensemble of storm scenarios based on local climatology, using a high-fidelity hydrodynamic model to compute surge responses for each scenario, and finally obtaining the statistics for the peak surge across different nodes within the geographic domain of interest. To reduce the burden in this hazard estimation (related to the high computational cost of the high-fidelity models) surrogate modeling is used. The choice of surrogate model is important because storm surge is governed by highly nonlinear relationship between storm characteristics and surge response patterns, compounded by complex spatial dependencies. This study investigates sparse graph-based surrogates, specifically graph neural networks (GNNs), for peak storm surge emulation. Sparsity is embedded into the graph connectivity, enabling the efficient application of GNNs to high-dimensional coastal domains while alleviating memory and scalability challenges. In addition, a response-based similarity metric is introduced to define graph edges, ensuring that the surrogate reflects the existing complex dependencies in surge responses across geo-locations. The results are evaluated in terms of predictive accuracy, computational efficiency, and memory requirements, providing a comprehensive assessment of the suitability of the proposed framework for storm surge em-

ulation.

Parisa Toofani Movaghar, Alexandros A. Taflanidis
University of Notre Dame
ptoofani@nd.edu, a.taflanidis@nd.edu

Norberto Nadal-Caraballo, Madison Yawn, Luce Aucoin
U.S Army Corps of Engineers
norberto.c.nadal-caraballo@usace.army.mil, madison.c.yawn@usace.army.mil, luke.a.aucoin@erdc.dren.mil

CP16

Conditional Sensitivity and Relative Uncertainty in Stochastic Oscillators Via Adjoint Methods

Quantifying the effect of stochastic perturbations on high-dimensional dynamical systems is a central challenge in uncertainty quantification. Oscillatory systems, such as neuronal networks or biochemical clocks, exhibit both intrinsic fluctuations and responses to external inputs, making it difficult to disentangle sources of variability. In this work, we show that the gradient of the stochastic asymptotic phase naturally provides a local, state-dependent sensitivity measure of phase to small perturbations. We interpret this gradient as a measure of relative uncertainty that separates the contribution of intrinsic noise from that of extrinsic inputs. This approach bridges classical adjoint sensitivity analysis, stochastic phase reduction, and conditional uncertainty quantification. Using examples of noisy limit-cycle systems, we demonstrate how this framework efficiently predicts phase diffusion, synchronization, and entrainment under additional perturbations, reducing high-dimensional stochastic dynamics to a low-dimensional phase description while preserving key uncertainty measures. Our results suggest that conditional adjoint sensitivities can serve as a general tool in UQ for characterizing the impact of perturbations relative to intrinsic stochastic variability, with potential applications in neuroscience, chemical oscillators, and networked stochastic systems.

Yangyang Du, Daniel Forger
University of Michigan
yyadu@umich.edu, forger@umich.edu

CP16

Sensitivity Analysis Using Stochastic Finite Volume Discretization with Fourier Amplitude Sensitivity Test

In this talk we present sensitivity analysis of partial differential equations using the Fourier Amplitude Sensitivity Test (FAST) in conjunction with the recently developed Stochastic Finite Volume (SFV) method. The SFV is a semi-intrusive method to quantify uncertainties from random inputs in partial differential equation (PDE) models, specifically hyperbolic PDEs. The FAST is an efficient global sensitivity technique that accounts for nonlinearity and non-monotonic relationships between model inputs and outputs. We will use evaluations pulled from the SFV in the FAST method to determine model input sensitivity for arbitrary quantities of interest. Our SFV+FAST technique proves computationally more efficient than traditional Monte Carlo methods for estimating sensitivities for complex models, which could be computationally prohibitive. We show the computation of the sensitivity using SFV+FAST for a series of models, ranging from idealized to real models based on experimental data.

Humberto C. Godinez, Chonggang Xu
Los Alamos National Laboratory
hgodinez@lanl.gov, cxu@lanl.gov

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

Jake J. Harmon
Los Alamos National Laboratory
harmon@lanl.gov

Nishant Panda
Los Alamos National Lab
nishpan@lanl.gov

Shengtai Li
Los Alamos National Laboratory
sli@lanl.gov

CP17

Bayesian Uncertainty Analysis of Prior Beliefs in History Matching Climate Models

Computer models are vital in climate science to study drivers of climate change; quantify uncertainties in future climate predictions; and to guide policy decisions. However, their direct use is inhibited by: their complex structure; high-dimensional input and output, including spatial-temporal fields; and their long evaluation times taking months on HPCs. Earth System Models (ESMs) integrate atmosphere, ocean, land, ice, and biosphere. ESMs are used to investigate natural and anthropogenic aerosol emission interactions with clouds yielding large Aerosol Radiative Forcing (ARF; temporal change in Earths energy balance due to aerosols) induced uncertainty in historical climate change. ARF is unobservable and key to predicting future climate yet research has resulted in little uncertainty reduction in 30-years of IPCC reports. Bayesian emulators are fast statistical approximations to model outputs. These are embedded in an uncertainty quantification framework incorporating structural model discrepancy linking ESM and the real-world, as well as representation and observation errors. This enables an efficient global parameter search for model-observation comparison via history matching to resolve parametric uncertainty and robustly constrain ARF for the UK Met Offices UKESM1 model. This requires many subjective choices: the output set used in history matching; the prior uncertainty specification; and the beliefs represented. We investigate the effect on ARF uncertainty.

Jonathan Owen
University of Sheffield
jonathan.owen@sheffield.ac.uk

CP17

Bayesian Variable Screening Via Joint Credible Region for Universal Kriging in Computer Experiments

Computer experiments using Gaussian process models are widely employed in scientific and engineering applications for emulating complex systems. Much of the existing work on variable selection for such models has focused on either mean function effects or covariance structure effects separately. This paper proposes a Bayesian variable screening approach that simultaneously evaluates variable impor-

tance in both components of Gaussian process models. A new joint credible region method is introduced to identify active variables. The key idea of the proposed method is to construct a two-dimensional credible region using posterior medians of an inert variable's regression coefficient and correlation parameter. This enables a unified decision framework of determining variable to be active or inactive based on whether their median posterior parameter pair lies inside or outside this joint region. The performance of the proposed method is evaluated through comprehensive simulation studies and real applications including the Borehole example and heat exchanger example.

Fan Zhang
Boise State University
fanzhang@boisestate.edu

Ray-Bing Chen
Department of Statistics
National Cheng Kung University
rbchen@ncku.edu.tw

Xinwei Deng
Virginia Tech
xdeng@vt.edu

MS1

Canonical Bayesian System Identification

Standard Bayesian approaches for linear time-invariant (LTI) system identification are hindered by parameter non-identifiability; the resulting complex, multi-modal posteriors make inference inefficient and impractical. We solve this problem by embedding canonical forms of LTI systems within the Bayesian framework. We rigorously establish that inference in these minimal parameterizations fully captures all invariant system dynamics (e.g., transfer functions, eigenvalues, predictive distributions of system outputs) while resolving identifiability. This approach unlocks the use of meaningful, structure-aware priors (e.g., enforcing stability via eigenvalues) and ensures conditions for a Bernstein–von Mises theorem—a link between Bayesian and frequentist large-sample asymptotics that is broken in standard forms. Extensive simulations with modern MCMC methods highlight advantages over standard parameterizations: canonical forms achieve higher computational efficiency, generate interpretable and well-behaved posteriors, and provide robust uncertainty estimates, particularly from limited data.

Andrey Bryutkin
Mathematics
MIT
bryutkin@mit.edu

Matthew Levine
Broad Institute
mattlevine22@gmail.com

Iñigo Urteaga
Basque Center for Applied Math
Bilbao, Spain
iurteaga@bcamath.org

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS1

Uncertainty Quantification in Synchronization Problems

Synchronization problemsrecovering group elements from noisy pairwise measurementsarise in computer vision, structural biology, and graph matching. A central challenge is not only to achieve accurate recovery but also to quantify uncertainty under severe corruption. I will present cycle-consistencybased methods, in particular Cycle-Edge Message Passing (CEMP), which provide both robust recovery and principled corruption estimates. CEMP offers corruption estimation and recovery guarantees for rotations and, more generally, for compact groups. These ideas extend naturally to partial permutations and to translations (without distance information), with applications to multiple image matching and camera pose estimation. Together, these results yield simple yet rigorous uncertainty scores on edgesboth provable and practically effective.

Yunpeng Shi
UC Davis
ypshi@ucdavis.edu

MS1

Data Augmentation: A Fourier Analysis Perspective

Data augmentation, which supplements a dataset with transformed copies of each datum according to a known symmetry group, provides a model-agnostic approach to enforcing invariances, in contrast to methods that encode symmetries directly into the model. Although data augmentation has proven effective in theory and practice, full group-sized augmentation is often computationally infeasible, prompting the question: Can sublinear augmentation still achieve the same performance as full augmentation in terms of generalization bounds and sample complexity? In this paper, we develop a theoretical framework based on Fourier analysis, showing that sublinear data augmentation can achieve the full statistical benefits achieved via full data augmentation. To our knowledge, this is the first proof of the efficacy of sublinear augmentation, highlighting an underexplored aspect of why augmentation remains a powerful and widely applicable strategy, even when performed only partially.

Melanie Weber
Harvard University
mweber@seas.harvard.edu

MS1

Statistical Learning Guarantees for Group-Invariant Barron Functions

We study the generalization error of group-invariant neural networks within the Barron framework and show that incorporating group-invariant structure leads to an approximation rate modulated by a group-dependent factor ($\delta_{G,\Gamma,\sigma} \leq 1$). When this factor is small, group invariance yields substantial improvements in approximation accuracy. On the estimation side, we establish that the Rademacher complexity of the group-invariant function class is no larger than that of its non-invariant counterpart, implying that the estimation error is not adversely affected by the inclusion of symmetry. As a result, the

overall generalization error can be significantly improved when learning target functions with inherent group symmetries. We further present illustrative examples demonstrating both favorable cases, where ($\delta_{G,\Gamma,\sigma} \approx |G|^{-1}$), and unfavorable cases, where ($\delta_{G,\Gamma,\sigma} \approx 1$). Overall, our results provide a rigorous theoretical foundation showing that encoding group-invariant structures in neural networks leads to clear statistical advantages for symmetric target functions.

Yahong Yang
Georgia Institute of Technology, U.S.
yyang3194@gatech.edu

MS2

Lifted Transition Operators for Fast Stochastic Physics Simulation

We introduce Lifted Transition Operators, which provides a method for fast simulation of stochastic systems. We combine high dimensional and low-dimensional stochastic states of the system to fit a transition operator that approximates the conditional distribution of a stochastic system. Our lifted transition operator achieves state-of-the-art accuracy on simulating trajectories of stochastic physics systems. Additionally, we show scaling to high complex tasks such as the generation of natural videos.

Jules Berman
NYU
jmb1174@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS2

On the Edge: Generative Model Assisted Sampling of Separtrices

A central challenge in modeling dynamical systems is the characterization of basins of attraction and the barriers, or separatrices, that divide them. Simulations in practice often become trapped in local minima over long timescales, and critical points of interest frequently lie along free energy barriers that are sparsely sampled. Reconstructing these separatrices is important not only for identifying tipping points of a system, but also for providing plausible data in regions of the phase space that are otherwise poorly populated. In parallel, advances in machine learning have introduced generative models that can learn to sample from complex target distributions and detect underlying low-dimensional structures. These models share conceptual parallels with physics-based enhanced sampling techniques, and can be used to sample densities from underlying manifolds present in systems of interest. In this work, we present a framework that integrates machine learning-based classification of basins of attraction with generative modeling to characterize, reconstruct, and visualize separatrices in dynamical systems, as well as quantify the uncertainty present in these steps.

Ellis R. Crabtree, Dimitris Giovanis, Ioannis G. Kevrekidis
Johns Hopkins University
ecrabtr2@jhu.edu, dgiovan1@jhu.edu, yannisk@jhu.edu

MS2

Witnessing Discrepancies in Imaging Modalities

Inspection of integrated circuits (advanced packaging) to identify defects and adversarial objects relies on a subjective examination of complicated images generated by X-ray computed tomography (xCT). The measurement is both time-intensive and human-dependent. We explore a method of visualizing and automatically quantifying discrepancies in random distributions of shapes given boundary segmentations of microchip xCT imaging. By motivating rigid invariant feature spaces as dual evaluation functionals informing a natural decomposition of segmented curves, we explore a kernel-based metric for comparing distributions of Separable Shape Tensors (SST's) over a learned product submanifold. These formalisms facilitate a framework testing a hypothesis that trusted images and a discovered images contain shapes from common distributions of generalized, anisotropic scale and nonlinear undulation. If we reject this hypothesis, subsequent witness functions can then be used to visualize subsets of shapes that correspond to significant discrepancies to inform rapid inspection and quantify differences.

Zach Grey
National Institute of Standards and Technology
Applied and Computational Mathematics Division
zachary.grey@nist.gov

MS2

Probabilistic Learning on Manifolds with Double Diffusion Maps

We present a generative learning framework for probabilistic sampling based on an extension of the Probabilistic Learning on Manifolds (PLoM) approach, which is designed to generate statistically consistent realizations of a random vector in a finite-dimensional Euclidean space, informed by a limited (yet representative) set of observations. In its original form, PLoM constructs a reduced-order probabilistic model by combining three main components: (a) kernel density estimation to approximate the underlying probability measure, (b) Diffusion Maps to uncover the intrinsic low-dimensional manifold structure, and (c) a reduced-order It Stochastic Differential Equation (ISDE) to sample from the learned distribution. A key challenge arises, however, when the number of available data points N is small and the dimensionality of the diffusion-map basis approaches N, resulting in overfitting and loss of generalization. To overcome this limitation, we propose an enabling extension that implements a synthesis of Double Diffusion Mapsa technique capable of capturing multiscale geometric features of the datawith Geometric Harmonics (GH), a nonparametric reconstruction method that allows smooth nonlinear interpolation in high-dimensional ambient spaces. This approach enables us to solve a full-order ISDE directly in the latent space, preserving the full dynamical complexity of the system, while leveraging its reduced geometric representation.

Dimitris Giovanis, Nikolaos Evangelou
Johns Hopkins University
dgiovan1@jhu.edu, nevange2@jh.edu.

Roger Ghanem
University of Southern California
Civil and Environmental Engineering

ghanem@usc.edu

Ioannis G. Kevrekidis
Johns Hopkins University
yannisk@jhu.edu

MS3

Weighted Quantization Using MMD: From Mean Field to Mean Shift Via Gradient Flows

Approximating a probability distribution using a set of particles is a fundamental problem in machine learning and statistics, with applications including clustering and quantization. Formally, we seek a weighted mixture of Dirac measures that best approximates the target distribution. While much existing work relies on the Wasserstein distance to quantify approximation errors, maximum mean discrepancy (MMD) has received comparatively less attention, especially when allowing for variable particle weights. We argue that a Wasserstein–Fisher–Rao gradient flow is well-suited for designing quantizations optimal under MMD. We show that a system of interacting particles satisfying a set of ODEs discretizes this flow. We further derive a new fixed-point algorithm called mean shift interacting particles (MSIP). We show that MSIP extends the classical mean shift algorithm, widely used for identifying modes in kernel density estimators. Moreover, we show that MSIP can be interpreted as preconditioned gradient descent and that it acts as a relaxation of Lloyd's algorithm for clustering. Our unification of gradient flows, mean shift, and MMD-optimal quantization yields algorithms that are more robust than state-of-the-art methods, as demonstrated via high-dimensional and multi-modal numerical experiments.

Ayoub Belhadji
Massachusetts Institute of Technology
abelhadj@mit.edu

MS3

Dynamical Sampling with Affine Invariant Gradient Flows and Beyond

The framework of gradient flows has played an important role in generating efficient dynamics for sampling probability distributions. I will present recent work on affine invariant gradient flows and their parametric versions for addressing sampling problems in high dimensions that are often anisotropic. This includes a new affine invariant exact ensemble MCMC sampler that achieves improved dimensional scaling, and an efficient approximate sampler in the context of variational inference with applications in Bayesian inverse problems. I will also briefly discuss a new affine invariant Hamiltonian sampler beyond gradient flows with state-of-the-art dimensional scaling.

Yifan Chen
UCLA
yifanchen@math.ucla.edu

MS3

Gaussian Particles Through Wasserstein Gradient Flow

Variational inference (VI) is a widely used framework in Bayesian analysis, where the posterior distribution is approximated within a chosen parametric family by minimizing a discrepancy measuretypically the KullbackLeibler (KL) divergence. In this work, we introduce a novel regularization of the VI objective using the Wasserstein distance over the variational family. We show that this regularization allows the variational inference problem to be reformulated as a projected Wasserstein gradient flow, which admits closed-form solutions when the parametric familly is a Gaussian distribution.We further show how the problem can be lifted to a Gaussian particle representation, allowing for efficient approximation of multimodal posterior distributions.

Marc Lambert
Inria
marc.lambert@inria.fr

MS3

HK Gradient Flows for Sampling: Two New Machine Learning Applications

This talk features two recent works on gradient flows for MCMC sampling in machine learning. The first part of this talk focuses on the global exponential decay of the entropy functionals along the unbalanced-transport gradient flows, a.k.a. Hellinger-Kantorovich (HK) or Wasserstein-Fisher-Rao. As an example, we will construct a new formal dissipation geometry, which we term the interaction-force transport, that guarantees convergence of minimizing the maximum-mean discrepancy. This then results in a new algorithm for sampling and generative modeling. In the second part, I will introduce a gradient flow formulation to enhance the robustness of learning and optimization algorithms under data distribution shifts. Previously, researchers have interpreted these problems as interacting gradient flows. In comparison, our new formulation unifies the analysis and motivates the design of new and flexible algorithms. In particular, I will show using entropy-regularized optimal transport for learning robustly under distribution shifts.

Jia-Jie Zhu
KTH Royal Institute of Technology
zplusj@gmail.com

MS4

Root-Finding for Fast Digital Twin Calibration

We pose the digital twin calibration problem as a simulation optimization where a sample-efficient stochastic metamodel is used to optimize for the best parameters. To accelerate the search, we propose to find the root of a signed discrepancy measure. We fit the metamodel on this signed metric and derive new acquisition functions and their gradients that empirically prove significant computational gains in a number of data-driven and physics-driven examples. We also analyze this framework when a root may not exist and provide error bounds for this approximate solution for calibration parameters.

Sara Shashaani, Yongseok Jeon
North Carolina State University
sshasha2@ncsu.edu, yjeon@ncsu.edu

MS4

Dynamic Parameter Calibration of Digital Twins Using Surrogate-Based Kalman Filter

Digital twinsvirtual representations of physical systems that evolve alongside their real-world counterpartsare transforming engineering, manufacturing, energy, and

healthcare. Digital twins complement physical assets by maintaining historical, current, and predicted states and behaviors. Unlike conventional simulations, digital twins are continuously updated using routinely collected sensor data to enable continuous monitoring, predictive maintenance, and informed decision-making. Their development often relies on physics-based first principles, with fidelity critically dependent on the accurate calibration of several parameters. We propose a dynamic calibration method that integrates a Kalman filter with a surrogate model, enabling real-time updates without costly simulator runs.

Eunshin Byon, Jaeshin Park
University of Michigan
ebyon@umich.edu, jaeshin@umich.edu

MS4

Validating a Stochastic Digital Twin with Multi-Epoch, Multi-Variate Data

We propose a nonparametric validation framework for a stochastic Digital Twin (DT) to verify if the distributions of the simulated and system-observed key performance indicators (KPIs) match. The system KPIs are non-stationary and non-identically distributed across epochs and only one realization is observed in each epoch rendering any validation method that requires independent and identically distributed data not applicable. We propose a hypothesis testing framework that consists of a set of marginal tests and one joint inter-correlation test. Each marginal test examines if the marginal distributions of the DT and system KPI within each epoch align in all observed epochs. For each epoch, the system KPI is converted via probability integral transform using the empirical marginal cumulative distribution function of the DT KPIs. The converted marginal KPIs from all epochs are then tested against a standard uniform distribution. If the marginal tests fail to reject, the joint test determines whether the inter-correlation of the DT and system KPIs, measured by their respective copula functions, match. We prove that our tests deliver a valid Type-I error upon appropriately controlling the simulation sample size with respect to the number of epochs. We further improve the finite-sample performance by adopting a step-down multiple hypothesis testing framework where a set of critical values are calculated using bootstrap.

Linyun He
HEC Montreal
linyun.he@hec.ca

Luke Rhodes-Leader
Lancaster University
l.rhodes-leader@lancaster.ac.uk

Eunhye Song
Georgia Institute of Technology
eunhye.song@isye.gatech.edu

MS4

Dynamic Calibration Framework for Digital Twins Using Active Learning and Conformal Prediction

This work presents an adaptive framework for dynamically calibrating digital twins (DTs) in response to evolving real-world (RW) conditions. Traditional simulation-based models often rely on fixed parameter estimates, limiting their adaptability over time. To address this, our approach integrates active learning (AL) with a dynamic calibration mechanism that keeps the DT aligned with RW observations. At each time step, a new data point is received, and a conformal prediction-based monitoring system assesses whether recalibration is needed. When a change in the RW system state is detected, DT parameters are updated using an efficient AL strategy. The framework reduces computational overhead by avoiding unnecessary DT evaluations while maintaining accurate system representation. We demonstrate the effectiveness of the proposed approach in achieving adaptive, cost-efficient DT calibration over time.

Ozge Surer
Department of Information Systems & Analytics
Miami University
surero@miamioh.edu

Xi Chen
Virginia Tech
xchen6@vt.edu

Sara Shashaani
North Carolina State University
sshasha2@ncsu.edu

MS5

Active Learning for Nonlinear Computer Model Calibration

Combining computer models with real observations is a powerful approach for analyzing complex physical systems. One of the key challenges in model calibration is accounting for uncertainty arising from model inadequacy. The KennedyOHagan (KOH) framework provides a principled way to address this issue when computer model outputs are linearly related to physical observations. In this work, we extend the KOH framework by introducing a nonlinear wrapping function that jointly models the relationship between design inputs and computer model outputs. We develop the corresponding statistical inference procedures and propose active learning strategies that simultaneously guide the collection of physical and computer model data. These strategies exploit closed-form expressions for the posterior mean and variance under specific kernel structures, leading to efficient and adaptive experimental design.

Andrews Boahen
Michigan State University
boahenan@msu.edu

MS6

Uncertainty Propagation in Intermittent Data Assimilation with Control-Based Importance Sampling

Intermittent data assimilation, whereby observations arrive at discrete-times, requires the propagation of a conditional distribution until the next observation occurrence. Incorrect or inaccurate approximation of the true initial distribution, modeling errors, stochastic forcing, and nonlinear dynamics combine to degrade the distribution approximation at the observation time. In this talk we focus on how to improve the uncertainty representation of the transition distribution in the finite sample case of particle filtering, and in light of the aforementioned difficulties. The central technique to achieve the improvement is the inclusion of a control-based technique to adjust the flow of the transition distribution, which can be understood as an active impor-

tance sampling strategy for the duration of the transition propagation and which provides samples from an optimal proposal. We discuss and illustrate with simple models the ability to counter the aforementioned difficulties and implicitly adjust for nonlinearities and account for model error.

Ryne Beeson
Princeton University
ryne@princeton.edu

MS6

Bridging the Model Hierarchy: A Physics-Guided Machine Learning Framework for High-Resolution Climate Simulation Enhancement

Operational models are high-resolution and contain many crucial variables. Despite numerous successes, biases persist in most of these models, particularly in identifying extreme events and accurately reproducing statistics. However, it is challenging to directly modify these models to improve their performance. Conceptual and intermediate coupled models accurately characterize certain features of nature. Yet, these coarse-grained models contain only a subset of variables within a specific domain. By leveraging the strengths of models of varying complexities, we develop a robust physics-driven machine learning modeling framework that bridges the model hierarchy through effective latent space data assimilation. The resulting model not only inherits operational model benefits, including high resolution and a comprehensive set of variables, but also obtains globally enhanced the accuracy through local dynamical and statistical improvements from simpler models. The latent space technique identifies the dominant nonlinear features of the underlying dynamics, which facilitates effective communication between models. Finally, the machine learning representation of the model significantly enhances simulation efficiency, providing high-quality datasets and advancing the quantification of uncertainty in extreme events. The framework has been applied to improve the performance of CMIP6 models in characterizing El Nio complexity by utilizing simpler yet statistically accurate models.

Pouria Behnoudfar
University of Wisconsin Madison
p.behnoud@gmail.com

Charlotte Moser
University of Wisconsin-Madison
crmoser2@wisc.edu

Sibo Cheng
CEREA, Ecole des ponts, Institut Polytechnique de Paris
sibo.cheng@enpc.fr

Nan Chen, Nan Chen
University of Wisconsin - Madison
chennan@math.wisc.edu, chennan@math.wisc.edu

MS6

Reduced Representations of Empirical Probability Distributions

We introduce a methodology for the reconstruction of empirical distributions from reduced representations of the same. The application we have in mind is in transmitting massive data sets when these consist of samples from high dimensional distributions, and our goal is to make the transmission process more efficient by requiring less data be used in the reconstruction. The strategy exploits the property of Bayesian inference and stochastic interpolation. It is assumed that the receiver of the reduced data set will have the requisite recovery map, which we formulate and compute. This talk will present the general methodology and encouraging preliminary results that Indicate that the strategy is viable and can be made practical on sizable, high dimensional distributions.

Juan M. Restrepo
Oak Ridge National Lab
juan.m.restrepo@gmail.com

Jorge Ramirez
ORNL
ramirezosojm@orn.gov

David M. Rogers
Oak Ridge National Laboratory
rogersdm@ornl.gov

Zezhong Zhang
Auburn University
zez0002@auburn.edu

MS6

Quantum Mechanical Closure of Partial Differential Equations with Symmetries

We develop a framework for the dynamical closure of spatiotemporal dynamics governed by partial differential equations. Employing the mathematical framework of quantum mechanics, we use the space of quantum states to model the unresolved degrees of freedom of the original dynamics and the method of quantum measurement to predict their contributions to the resolved dynamics. The embedded dynamics is projected to finite dimension by a positivity preserving discretization process. Relying on kernel methods and time delay embedding, the compressed finite dimensional representation of the dynamics is invariant under the spatial symmetries of the original dynamics. We apply the data driven formulation of our scheme to closure problems for the shallow water and Kuramoto-Sivashinsky equations, demonstrating that it can accurately predict the main features of the true dynamics for out of sample initial conditions.

Chris Vales, David Freeman
Dartmouth College
chris.vales@dartmouth.edu,
david.c.freeman.gr@dartmouth.edu

Joanna M. Slawinska
Department of Mathematics
Dartmouth College
joanna.m.slawinska@dartmouth.edu

Dimitrios Giannakis
Dartmouth College
Dimitrios.Giannakis@dartmouth.edu

MS7

Intrepid MCMC Sampler for Structural Reliability Estimation Via Subset Simulation

Monte Carlo simulation augmented with sampling variance reduction schemes are among the most widely adopted

techniques in computational reliability (failure probability) estimation. When limit state functions (LSFs) possess geometric features such as (a) discontinuities or sharp nonlinearities, or (b) several, possibly disconnected regions of failure, even advanced variance reduction techniques, such as the subset simulation method, may fail to deliver accurate estimates of failure probability. In such cases, the underlying Markov chain Monte Carlo (MCMC) sampler may fail to detect and adequately sample from all modes of the conditional pdf that contribute to the probability of failure. This study introduces an improved MCMC sampler, termed Intrepid MCMC, which explores the parameter space more effectively, leading to improved sampling from important failure regions. The key idea is to utilize two Markov transition pdfs, (a) one focusing on sampling from the mode of the conditional pdf in which the Markov chain currently resides, and (b) another dedicated to exploring the parameter space in search of previously undetected modes. These transition pdfs are combined to yield a composite transition pdf which leads to improved sampling from the target conditional pdf. Intrepid MCMC has previously shown to significantly outperform conventional random walk Metropolis methods for multimodal distributions, and we now demonstrate its effectiveness for complex reliability problems.

Adwait Sharma, Promit Chakroborty, Michael D. Shields
Johns Hopkins University
ashar125@jh.edu, pchakro1@jhu.edu,
michael.shields@jhu.edu

MS7

Active Learning-based Surrogate Modeling Method for Efficient Reliability Analysis of Stochastic Structural Dynamical Systems

It is important to evaluate the reliability of structural dynamical systems under various uncertainties. Surrogate models are widely employed to enable efficient reliability analysis in complex, computationally intensive engineering problems. However, building high-accuracy surrogate models for dynamic reliability estimation with limited computational resources remains a significant challenge. This work presents a novel active learning-based surrogate modeling method for the efficient estimation of first-passage failure probabilities of stochastic structural dynamical systems. In this approach, Kriging surrogate models are constructed in a latent functional space obtained through functional dimension reduction. By leveraging the prediction uncertainty and the concept of misclassification probability, a new learning function is then developed to guide the selection of the best next sample for model enhancement. Furthermore, an error-based stopping criterion is proposed to judge the convergence of the active learning process. Through numerical examples of varying dimensionality and complexity, it is shown that the proposed method is efficient and accurate for dynamic reliability analysis.

Zhouzhou Song
TU Dortmund University
zhouzhou.song@tu-dortmund.de

Chao Dang, Marcos Valdebenito, Matthias Faes
Technical University of Dortmund
chao.dang@tu-dortmund.de,
mailto:marcos.valdebenito@tu-dortmund.de,
matthias.faes@tu-dortmund.de

MS7

Mathematical Analysis of Variance-Based Reliability Sensitivity Indices

In this work we analyze the sensitivity of indicator random variables associated with a failure event that occurs with a small probability. This is a classical problem in engineering risk and reliability analysis. For the sensitivity analysis we use the well known Sobol' indices defined in terms of the variance decomposition of the failure indicator random variable. We derive tight two-sided bounds for the closed and total Sobol' indices of affine linear limit-state functions and standard Gaussian inputs. In addition we show that the closed and total Sobol' indices approach either zero or one in the limit of the failure probability going to zero. This proves rigorously an empirical observation in previous studies. Moreover, we study second-order approximations of reliability sensitivity indices and show their asymptotic behavior under certain conditions. We illustrate the theoretical results by a simple test problem with an affine linear limit-state function. Moreover, we consider a practical engineering test problem with a nonlinear limit-state function.

Elisabeth Ullmann
Technical University of Munich
elisabeth.ullmann@tum.de

Iason Papaioannou
Engineering Risk Analysis Group
Technical University of Munich
iason.papaioannou@tum.de

MS7

Failure Probability Estimation Using Physics-Based Surrogate Models

In this talk, we present a failure probability estimation approach based on domain-specific surrogate models. The key hypothesis is that computationally efficient, simplified models often exist for complex scientific applications. By leveraging the failure domain of such a surrogate model, we construct a conditional sampling scheme to approximate the failure probability of the more complex model. The proposed method scales well with dimensionality and applies to first-passage problems involving random fields and stochastic processes.

Ziqi Wang
University of California, Berkeley
ziqiwang@berkeley.edu

MS8

Efficient Multilevel Methods for Stochastic Material Properties Inversion in Heat Transfer Problems

In this work, we extend multilevel methods - originally developed for linear-quadratic distributed optimal control of elliptic equations - to the problem of identifying a stochastic diffusivity model in heat transfer from state measurements. Using synthetic experiments with applied heat sources and multiple measurements of temperature and heat flux, our goal is to recover the diffusivity tensor. For this study, the tensor is assumed isotropic and represented by a truncated KarhunenLove expansion. Because the problem is ill-posed, we formulate it as a regularized least-

squares optimization, which is nonlinear and non-convex. To solve it, we use a hierarchy of geometric grids and multilevel techniques to precondition the Hessian in the NewtonCG process.

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

Liviu Marin
Department of Mathematics
University of Bucharest, Romania
liviu.marin@fmi.unibuc.ro

Alex Sheranko
Department of Mathematics and Statistics
University of Maryland, Baltimore County
asheran1@umbc.edu

MS8

Nonlinear Reduction Model for Solving Parameterized PDE Problems Using Machine Learning Techniques

In many engineering systems, solving parameterized partial differential equations (PDEs) over many parameter values is inevitable. Reduced order modelling (ROM) replaces the full-order model by a model with smaller dimension. Linear ROMs, e.g., reduced basis methods often perform poorly when the Kolmogorov n-width of a solution manifold decays slowly. To address this, we propose a new method that consists of two stages: first, finding a low-dimensional space and second, use a feedforward neural network to map this space to high-fidelity solutions. Thus, our surrogate model uses two nonlinear mappings. The first mapping is parameter-independent and maps from the reduced basis space to a low-dimensional space. The second mapping is parameter-dependent and maps the low-dimensional space to the solution space. We discuss the above solution strategy and apply it to some parameterized PDEs for which linear ROMs do not perform well. In the first test, a 1D Burgers' equation with solutions exhibit strong oscillations provides a difficult example. The second test is a 2D lid-driven cavity problem over a high dimensional parameter domain. Results show that our approach improves the accuracy compared to linear ROMs. Furthermore, the nonlinear mapping that constructs a low-dimensional representation also improves the training efficiency in the second stage. These advantages make our strategy suitable for fast simulations of parameterized PDEs requiring multiple solves.

Jingye Li, Jinglai Li
University of Birmingham
jxl1728@student.bham.ac.uk, j.li.10@bham.ac.uk

Alex Bespalov
University of Birmingham, United Kingdom
a.bespalov@bham.ac.uk

MS8

Multigrid Preconditioners for Optimal Control Problems with Stochastic Elliptic PDE Constraints

In this work we construct multigrid preconditioners to be used in the solution process of pathwise optimal control problems constrained by elliptic partial differential equations with random coefficients. We use a sparse-grid collocation method to discretize in the stochastic space and multigrid techniques in the physical space. Numerical results show that the proposed preconditioners lead to significant computational savings, with the number of preconditioned conjugate gradient iterations decreasing as the resolution in the physical space increases.

Ana Maria Soane
U.S. Naval Academy
Mathematics Department
soane@usna.edu

MS8

Learning-Based Algorithm for Eigenvalue Problems with Parametric Uncertainty

We study eigenvalue problems under parametric uncertainty. The matrix operator is assumed to be given in the form of a real affine expansion. We propose several computational strategies that increase efficiency of describing several eigenvalues of interest, including their possible paths and/or probability distributions by Monte Carlo simulation or spectral stochastic finite elements. To this end we also utilize a neural-network surrogate trained by a low-fidelity simulation to provide an estimate of the eigenvalue path in the complex plane. The effectiveness of the method is illustrated by numerical experiments using a benchmark problem and also a linear stability assessment of incompressible NavierStokes equations with stochastic viscosity.

Bedrich Sousedik
University of Maryland Baltimore County
sousedik@umbc.edu

Ahmet Kaan Aydin
University of Maryland - Baltimore County
aaydin@umbc.edu

MS9

Multi-Fidelity Uncertainty Quantification of Dynamical Systems Using Structure-Preserving Operator-Inference Neural Networks

We present an approach to operator inference using neural network (NN) models and demonstrate it in the context of multi-fidelity uncertainty quantification (UQ). We infer a composite NN operator comprised of a series of structure-preserving terms e.g., matrices with symmetry and definiteness properties from state observation data. As opposed to classic operator inference approaches, the use of an NN model class enables representation of generic non-linearities at the expense of non-convex inference. We demonstrate that this approach provides superior accuracy over several existing NN-based inference procedures and demonstrate its ability to reduce variance in a forward-propagation UQ problem with uncertain parameters.

Eric Parish
Sandia National Lab
eric.parish9@gmail.com; ejparis@sandia.gov

Anthony Gruber
Sandia National Laboratories
adgrube@sandia.gov

MS9

Entropy Stable Reduced Order Modeling of Nonlinear Conservations Laws Using Entropy Correc-

tion Artificial Viscosity

Reduced order models (ROMs) are inexpensive surrogate models that reduce costs associated with many-query scenarios, but face challenges of robustness and stability when applied to nonlinear conservation laws. These can be addressed by constructing ROMs which are "entropy stable"; however, existing constructions rely on relatively invasive and expensive flux differencing formulations. In this work, we propose an alternative method of constructing an entropy stable ROM based on "entropy correction artificial viscosity" (ECAV), which adds a minimally dissipative correction to an arbitrary ROM to enforce a semi-discrete entropy inequality. We also propose a conservative hyper-reduction technique for finite volume FOMs, which, when combined with ECAV, results in a conservative and entropy stable ROM.

Ray Qu, Jesse Chan
Rice University
rq10@rice.edu, jesse.chan@rice.edu

MS9

Preconditioning for Adjoint Sensitivity Analysis with Application to Optimization of Radiation Transport

Achieving robust control and optimization in high-fidelity physics simulations is extremely challenging, especially for evolutionary systems whose solutions span vast scales across space, time, and physical variables. In conjunction with gradient-based methods, adjoint systems are widely used in the optimization of systems subject to differential equation constraints. In optimization, gradient-based methods are often transformed using suitable preconditioners to stabilize and accelerate the convergence of the optimization algorithm. Inspired by preconditioned gradient descent methods, we introduce a framework for the preconditioning of adjoint systems associated to evolution equations, which allows one to reshape the dynamics of the adjoint system. Using techniques from symplectic geometry and Hamiltonian mechanics, we show that these preconditioned adjoint systems preserve the property that the adjoint system backpropagates the derivative of an objective function. We then apply this framework to the setting of coupled evolution equations, where we develop a notion of scale preconditioning of the adjoint equations when the state dynamics exhibit large scale-separation. We demonstrate the proposed scale preconditioning on an inverse problem for the radiation diffusion equations.

Brian K. Tran
University of Colorado Boulder
Department of Applied Mathematics
brian.tran@colorado.edu

MS9

Structure Preserving Data Driven Models from Experiment with Quantified Uncertainty

TBD

Nat Trask
University of Pennsylvania, U.S.
ntrask@seas.upenn.edu

MS10

Rare Event Estimation Using Tensor Trains and Large Deviation Theory

Recently, algorithms for approximating and sampling from arbitrary multivariate probability densities using low-rank tensor-based surrogates have been proposed. In particular, the deep inverse Rosenblatt transport (DIRT) algorithm, which uses a layered composition of functional tensor train approximations, has found success in a range of uncertainty quantification tasks. In this work, we propose some practical strategies to accelerate the application of the DIRT algorithm to rare event estimation problems. When the uncertainty in the system of interest is characterised by a complex distribution without a known analytical form (e.g., a Bayesian posterior), we use the DIRT algorithm to construct a coupling between this distribution and the standard Gaussian distribution. This allows us to apply ideas from large deviation theory to design preconditioning techniques to efficiently construct a DIRT approximation to the optimal importance sampling distribution for the rare event probability. We illustrate our methodology using several numerical examples.

Alex De Beer
University of Sydney
adeb0907@uni.sydney.edu.au

Karina Koval
Heidelberg University
karina.koval@iwr.uni-heidelberg.de

Tiangang Cui
Monash University
tiangang.cui@sydney.edu.au

MS10

Low-Rank Approximations for Large-Scale Nonlinear Feedback Control

Computation of the optimal feedback law for general (nonlinear/unstable/stochastic) dynamical systems requires solving the Hamilton-Jacobi-Bellman Partial Differential Equation (PDE), which suffers from the curse of dimensionality. We develop a unified framework for computing a fast surrogate model of the feedback control function based on low-rank decompositions of matrices and tensors. Firstly, we propose a Statistical Proper Orthogonal Decomposition (SPOD) for Model Order Reduction of very high-dimensional systems, such as the discretized Navier-Stokes equation or other PDEs, by compressing snapshots corresponding to random samples of all parameters in the system, initial condition and time. Secondly, we compute a low-rank Functional Tensor Train (TT) approximation of the feedback control function for the reduced model. Thus pre-trained TT representation of the control function of the reduced state can be used for real-time online generation of the control signal. Using the proposed SPOD and TT approximations, we demonstrate a controller computable in milliseconds that achieves lower vorticity of the Navier-Stokes flow with random inflow compared to using the mean inflow to produce reduced bases or controllers.

Sergey Dolgov
University of Bath
sd901@bath.ac.uk

Dante Kalise
Imperial College London
d.kalise-balza@imperial.ac.uk

Luca Saluzzi
Sapienza
University of Rome
luca.saluzzi@uniroma1.it

MS11

Novel Optimal Criteria in Bayesian Optimal Experimental Design

Bayesian optimal experimental design (OED) provides a principled framework for choosing experimental settings that maximize the informativeness of collected data. Traditional Bayesian OED has relied heavily on the expected information gain (EIG), based on KullbackLeibler divergence, or variance-reduction criteria such as A-optimality. In this talk, we introduce a new design criterion grounded in optimal transport: the expected Wasserstein distance between prior and posterior distributions. This criterion inherits the fully Bayesian character of EIG while offering geometric and stability advantages. For Gaussian regression models, the Wasserstein-2 criterion yields explicit formulas akin to classical optimality measures, making it computationally tractable. A central contribution of our work is a stability analysis of the Wasserstein-1 criterion, demonstrating robustness under perturbations of the prior and likelihood and enabling rigorous error control when empirical priors are used.

Tapio Helin
LUT University
Tapio.Helin@lut.fi

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

Rodrigo Rojo-Garcia
LUT University
Finland
rodrojog89@gmail.com

MS11

Bayesian Experimental Design for Prediction Using Transport Approaches

We consider Bayesian optimal experimental design (BOED) with a focus on maximizing expected information gain (EIG), a task that is often computationally challenging due to the dependence of the EIG on intractable densities. In many settings, samples from the joint distribution of parameters and data are available, and these can be used to construct transport-map-based approximations to intractable densities (e.g., posteriors, evidence, or priors) by minimizing a suitable loss function. We present a sample-based transport approach for approximating the EIG in such settings. The methodology is illustrated through numerical examples.

Karina Koval
Heidelberg University
karina.koval@iwr.uni-heidelberg.de

Tiangang Cui
Monash University
tiangang.cui@sydney.edu.au

MS11

Optimal Design of Bayesian Inverse Problems with Mobile Sensor Paths

We study optimal experimental design (OED) for inverse problems governed by PDEs, where the inversion parameter lies in an infinite-dimensional Hilbert space. In our framework, we parameterize sensor trajectories and treat the path parameter as the experimental design. This yields a general formulation compatible with standard OED criteria, including the c-optimality criterion. The problem is discretized with finite elements and solved with gradient-based methods. Numerical experiments on an advection-diffusion equation demonstrate the approach: we solve OED problems with two different path parameterizations and show large improvements over random or user-selected paths. Our framework is general and extensible to mobile sensing in real applications, with relevance to wildfire prediction and control.

John N. Neuberger
North Carolina State University
jnneuber@ncsu.edu

Bart Waanders
Sandia National Labs
bartv@sandia.gov

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

MS11

Sequential Bayesian Optimal Experimental Design in Infinite Dimensions via Policy Gradient Reinforcement Learning

Sequential Bayesian optimal experimental design (SBOED) for PDE-governed inverse problems is computationally demanding for infinite-dimensional random field parameters, due to repeated forward and adjoint solves within nested Bayesian inversion. We formulate SBOED as a finite-horizon Markov decision process and learn an amortized design policy using policy-gradient reinforcement learning (PGRL), enabling online design selection without repeatedly solving optimization problems. We combine active subspace projection for parameters and principal component analysis for states with a derivative-informed latent attention neural operator (LANO) surrogate that predicts both the parameter-to-solution map and its Jacobian. We adopt a Laplace-based D-optimality reward and introduce an eigenvalue-based evaluation strategy that uses prior samples as proxies for MAP points. Numerical experiments on sequential multi-sensor placement for contaminant source tracking show approximately 100× speedup over high-fidelity finite element methods, improved performance over random placement, and interpretable upstream tracking policies.

Kaichen Shen, Peng Chen
Georgia Institute of Technology
kshen77@gatech.edu, pchen402@gatech.edu

MS12

Fokker-Planck-Inverse Reinforcement Learning: a Physics-Constrained Approach to Markov Decision

Process Models of Cell Dynamics

Inverse Reinforcement Learning (IRL) seeks to estimate the unknown reward function of a Markov Decision Process (MDP) from observed agent trajectories. IRL needs a transition function, and most algorithms assume it is known or can be estimated from data. It therefore becomes challenging when such transition dynamics is not known a-priori, since it enters the estimation of the policy. When the dynamics of agents in the state-action space is described by stochastic differential equations (SDE) in Itô calculus, these transitions can be inferred from the associated Fokker-Planck (FP) equation. We conjecture an equivalence between the time-discrete FP and MDP. We identify specific manifestations of this equivalence and use them to create a novel physics-aware IRL algorithm, FP-IRL, which can simultaneously infer the transition and reward functions using only observed trajectories. We employ PDE inference for the potential function in FP, which consequently allows the evaluation of reward, transition, and policy by leveraging the conjecture. We demonstrate the effectiveness of FP-IRL by applying it to a synthetic benchmark and a biological problem of cancer cell dynamics, where the transition function is inaccessible.

Chengyang Huang
University of Michigan
chengyah@umich.edu

Siddhartha Srivastava
Auburn University
sis0024@auburn.edu

Kenneth Ho, Katherine Luker, Gary Luker, Xun Huan
University of Michigan
kennetho@umich.edu, kluker@med.umich.edu, gluker@umich.edu, xhuan@umich.edu

Krishna Garikipati
USC
garikipa@usc.edu

MS12

Robust Learning of Long-time Macroscopic Dynamics from Short-time Particle Simulations with Efficient Uncertainty Quantification

Dissipative phenomena, ranging from diffusion in gases to plasticity in crystalline materials, is critical to many technological sectors. Yet, their continuum modeling necessarily requires strong phenomenological assumptions and their direct particle simulations at macroscopic length- and timescales remain computationally prohibitive. While machine learning offers promising avenues to develop constitutive models and learn PDEs from data, existing approaches often struggle with issues like the non-uniqueness of the underlying thermodynamic structure for given macroscopic dynamics and the lack of efficient uncertainty quantification. Here, we present SPIEDiff (Statistical-Physics Informed Epistemic Diffusion Models), a machine learning framework designed to coarse-grain particle dynamics while overcoming the aforementioned issues. SPIEDiff leverages fluctuation-dissipation relations from statistical physics to learn the unique dissipative operator guiding the dynamics, diffusion models, to efficiently deal with noisy and scarce data, and epinets, to quantify the epistemic uncertainty. SPIEDiff yields fast and accurate long-time continuum predictions for arbitrary boundary conditions from short-time particle simulation data, and it exhibits superior performance and robustness compared to the prior baseline models, Stat-PINNs (Statistical-Physics Informed Neural Networks) and its variant for uncertainty quantification.

Zequn He
University of Pennsylvania
hezequn@seas.upenn.edu

Celia Reina
University of Pennsylvania
Penn Institute for Computational Science
creina@seas.upenn.edu

MS12

Multilevel Bayesian Modeling of Local Growth Probabilities in Intracranial Aneurysms

Cerebral aneurysms cause nearly half million deaths annually worldwide, yet rupture risk assessment remains subjective. Monitoring relies on manual 2D MRA annotations, inconsistent across operators and insensitive to subtle 3D changes. We propose a Bayesian framework that quantifies growth from longitudinal 3D segmentations of aneurysm and vessel. Vessel walls are assumed rigid with zero-mean displacement over time, and their distribution captures uncertainty from resolution, segmentation, and registration. At the global level, aneurysm displacement distribution are compared against the vessel distribution using Wasserstein distance. The distance is then modeled with a Bayesian soft threshold, turning noisy measurements into global growth probabilities with posterior uncertainty. This yields interpretable cut-offs, slopes, and error terms that update as data accumulate. At the patient level, we extend growth probabilities locally to surface vertices. Rigid vessel vertices define the uncertainty model, while aneurysm vertices follow patient-specific lognormal growth distributions. The global probability acts as a prior for each vertex, producing maps that highlight high-risk regions. These maps guide borderline cases, reveal growth hotspots, and link with hemodynamic metrics for biomarker discovery. By combining global thresholds with local maps, our multilevel Bayesian model delivers risk scores and surgical guidance, reducing subjectivity and improving confidence.

Abhishek Singh, Atharva Hans, Jorge Roa Castro, Kostiantyn kondratiuk
Purdue University
sing1062@purdue.edu, hans1@purdue.edu, jroacast@purdue.edu, kkondra@purdue.edu

David Saloner
University of California, San Francisco
david.saloner@ucsf.edu

Vitaliy L Rayz, Ilias Bilionis, Pavlos Vlachos
Purdue University
vrayz@purdue.edu, ibilion@purdue.edu, pvlachos@purdue.edu

MS13

Quantifying Unparameterized Spatial Uncertainties in Deep Geologic Repositories for Nuclear Waste

Spatially heterogeneous uncertainties are prevalent in geophysical modeling applications, such as probabilistic post-closure performance assessment (PA) of deep geologic

repositories for nuclear waste. Such uncertainties are often highly influential to model outputs, so it is desirable to identify the most important mechanisms by which they influence model predictions. However, these uncertainties are often unparameterized in the sense that there is no set of parameters that can be specified to yield a particular realization of the uncertainty. Additionally, the uncertainty is not intrinsically endowed with a parameterization that captures a realizations mechanistic influence on model outputs. Therefore, in this talk we present a novel methodology to develop and assess a set of proxy variables that aim to represent this influence. We show how they can be computed, downselected, and incorporated into the construction of statistical surrogate models mapping model inputs to outputs. We present our methodology in the context of a motivating application problem in deep geologic repository PA and discuss the challenges in capturing the effects of these spatial heterogeneities in uncertainty analyses. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Teresa Portone, Dusty Brooks
Sandia National Laboratories
tporton@sandia.gov, dbrooks@sandia.gov

Laura P. Swiler
Retired (Sandia National Laboratories)
lpswiler@msn.com

MS13

Kernel Density Estimation Adaptive Importance Sampling with Gaussian Processes for Rare-Events Probability Estimation

Estimating rare–event probabilities for expensive black–box models is notoriously difficult because failures are scarce, structured, and a priori unknown. We present an adaptive importance sampling strategy that learns a proposal nonparametrically from data via a kernel density surrogate, allowing it to track multimodal and anisotropic failure regions without committing to a parametric family. The method interfaces with user–defined input laws and simple domain constraints, and it is designed to be plug–and–play within existing UQ workflows. From a theoretical standpoint, we show that the resulting estimator is unbiased at every iteration and consistent as the design adapts. Under mild regularity conditions, we establish a central limit theorem with an explicit asymptotic variance that contracts toward the oracle variance associated with the ideal failuretilted density. We also provide finitesample diagnostics and stability guarantees that prevent weight degeneracy during adaptation. Across reliability benchmarks, the approach delivers substantial variance reduction relative to static IS while maintaining conceptual and implementational simplicity.

Ashwin Renganathan
Pennsylvania State University
ashwin.renganathan@psu.edu

Annie Booth
Virginia Tech
annie_booth@vt.edu

MS13

Fourier Neural Operator-based Emulators for Coupling Engineering and Environmental System Simulations: Evaluating the Environmental Impact of Advanced Nuclear Reactor Spent Fuel

Tackling environmental challenges requires the understanding and models in both engineering systems in which waste and toxic substances are created and environmental systems in which contaminants can migrate and affect the public health. The integrated analysis of both engineering and environmental systems is critical for quantifying the actual environmental impact and health consequences. At the same time, the integration allows us to identify the engineering and design parameters and other factors that can potentially reduce and minimize the environmental impact. In this study, we develop an engineering/environmental system integration framework through deep-learning (DL)-based emulators, particularly using the Fourier Neural Operators (FNO). Each model and emulator are developed separately for an engineering system and an environmental assessment model, and then fast running emulators enable the rapid integration. We demonstrate our approach to evaluate multiple advanced nuclear reactors from the environmental impact of their spent nuclear fuel (SNF). This integrated system analysis will quantify the impact of different reactor/fuel designs on the waste repository performance and potential health risk as well as optimize reactor/processing designs for minimizing the environmental impact.

Haruko M. Wainwright
Massachusetts Institute of Technology
hmwainw@mit.edu

MS13

Scalable Gaussian Processes with Low-rank Deep Kernel Decomposition

Kernels are key to encoding prior beliefs and data structures in Gaussian process (GP) models. The design of expressive and scalable kernels has garnered significant research attention. Deep kernel learning enhances kernel flexibility by feeding inputs through a neural network before applying a standard parametric form. However, this approach remains limited by the choice of base kernels, inherits high inference costs, and often demands sparse approximations. Drawing on Mercer's theorem, we introduce a fully data-driven, scalable deep kernel representation where a neural network directly represents a low-rank kernel through a small set of basis functions. This construction enables highly efficient exact GP inference in linear time and memory without invoking inducing points. It also supports scalable mini-batch training based on a principled variational inference framework. We further propose a simple variance correction procedure to guard against overconfidence in uncertainty estimates. Experiments on synthetic and real-world data demonstrate the advantages of our deep kernel GP in terms of predictive accuracy, uncertainty quantification, and computational efficiency.

Yao Xie
Georgia Institute of Technology
yao.xie@isye.gatech.edu

MS14

Earthquake Forecasting with Uncertainty Quantification

Earthquake forecasting is critical for hazard assessment but remains extremely challenging due to unobservable state variables, sparse and low-resolution observations, and the multiscale dynamics of the earthquake cycle. To address

these difficulties, we develop machine-learned reduced-order models (ROMs) of physics-based earthquake simulations and use them within ensemble-based data assimilation. With the Ensemble Kalman Filter (EnKF), we estimate the evolving system states from partial observations, enabling probabilistic forecasts of the timing and spatial distribution of future events together with uncertainty quantification. In a related direction, we apply Ensemble Kalman Inversion (EnKI) to infer uncertain physical parameters from sparse paleoseismic records, obtaining both parameter estimates and their uncertainties. These efforts illustrate how reduced-order modeling, data assimilation, and inversion methods can tackle the dual challenges of forecasting earthquake events and constraining the governing physical models.

Hojjat Kaveh
Caltech
hkaveh@caltech.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

Jean Philippe Avouac
Caltech
avouac@caltech.edu

MS14

Accelerating Uncertainty Quantification of Magnetotelluric Data Using Neural Operators

The magnetotelluric (MT) method is a passive source geophysical method to image the conductivity distribution of the subsurface. This is achieved by measuring the natural electric field and the magnetic field at the earth's surface and then obtaining the transfer function between the two measurements. The imaging problem involves obtaining the conductivity distribution that can reproduce the transfer functions. It is an ill-posed and underdetermined system solved with the aid of a regularizer. Markov Chain Monte Carlo (MCMC) methods have been used widely to quantify the uncertainties in the model estimates obtained from the inverse problem. This requires solving the governing differential equations multiple times. The computation, usually done through traditional numerical methods is expensive, limiting their usage to 1D. Neural operators can be used to approximate the differential equation operator, in our case the operator mapping the conductivity distribution to the electric and the magnetic field. Since the regularizer informs the *a priori* distribution, usually one consisting of smooth models, limiting ourselves to this family guarantees the accuracy of learned surrogates while providing interpretable subsurface features. We validate our results against the ones obtained using MCMC and show that the obtained surrogate, being orders of magnitude faster than the traditional methods, enables us to perform uncertainty quantification efficiently.

Abhinav Pratap Singh, Samer Naif
Georgia Institute of Technology
abhinavps@gatech.edu, snaif@eas.gatech.edu

MS14

Randomized Blocky Occam for Generating 2D Resistivity Models with Sharp Contrasts from Electromagnetic Data

Geological boundaries such as sedimentary-volcanic contacts, groundwater tables, magmatic bodies, and faults produce sharp resistivity contrasts critical for subsurface characterization. Standard electromagnetic (EM) inversion methods generate smooth models that blur these features. Existing approaches for generating sharp resistivity contrasts are computationally expensive and difficult to use. We present Blocky Occam, a new inversion method that enforces sharp resistivity contrasts using total variation (TV) regularization. TV regularization has been applied successfully in compressed sensing and image deblurring, and we show how to use it efficiently in EM inversion. We integrate TV regularization within the established Occams inversion: Linearize the model, solve a sequence of linear TV regularized problems efficiently via split Bregman to determine an appropriate regularization, and iterate this process until convergence. Randomized blocky Occam (RamBO) adapts the randomize-then-optimize (RTO) for uncertainty quantification using edge-preserving Laplace priors. We demonstrate the robust efficiency of blocky Occam and RamBO on synthetic and real (field data) 2D EM inverse problems. Our results indicate that RamBO improves the resolution of geological boundaries and produces reliable uncertainties. The algorithm is computationally efficient, parallelizable, and enables practical application of blocky inversions to complex geological settings.

Eliana Vargas Huitzil
Scripps Institution of Oceanography
evargashuitzil@ucsd.edu

Matthias Morzfeld
University of California, San Diego
Scripps Institution of Oceanography
mmorzfeld@ucsd.edu

Steven Constable
Scripps Institution of Oceanography
University of California, San Diego
sconstable@ucsd.edu

MS15

Ensemble Kalman Updates: a Non-Asymptotic Approach

Filtering is concerned with online estimation of the state of a dynamical system from partial and noisy observations. In applications where the state of the system is high dimensional, ensemble Kalman filters are often the method of choice. These algorithms rely on an ensemble of interacting particles to sequentially estimate the state as new observations become available. Despite the practical success of ensemble Kalman filters, theoretical understanding is hindered by the intricate dependence structure of the interacting particles. This talk will discuss ensemble Kalman filters that incorporate an additional resampling step to break the dependency between particles. The new algorithm is amenable to a non-asymptotic and dimension-free theoretical analysis that extends and improves upon those available for filters without resampling, while also performing well in numerical examples. Papers: https://arxiv.org/abs/2208.03246 https://arxiv.org/pdf/2308.08751.pdf

Omar H. Al-Ghattas
Broad Institute

oag@broadinstitute.org

MS15

A Bayesian Characterization of Ensemble Kalman Updates

The Ensemble Kalman Update (EnKU), an update in the Ensemble Kalman Filter (EnKF), is widely used for Bayesian inference in inverse problems and data assimilation. At each step, it solves a likelihood-free Bayesian inversion problem using an ensemble of particles drawn from a joint measure and a single observation. The empirical posterior is formed by transporting the particles through a linear map given by the Kalman gain. Although EnKU is exact for Gaussian joints in the mean-field, exactness alone does not determine a unique update. Many affine maps send a Gaussian joint to the same posterior, so which affine map should we use? We give a new characterization of EnKU among all such maps. First, we identify the set of laws for which EnKU yields exact conditioning and show that it is much larger than the Gaussian family. Next, we prove that, except for a small class of highly symmetric distributions (including Gaussians), EnKU is the unique exact affine conditioning map. Finally, we ask for the largest possible collection of measures for which any measure-dependent affine transport could be exact. We show that EnKUs exactness set is nearly maximal, differing only by a small symmetry class. Thus, among affine transports, EnKU is near-optimal for exact distributional conditioning beyond the Gaussian case, and it remains uniquely exact except for a narrow subclass of highly symmetric distributions.

Frederic Jorgensen
Massachusetts Institute of Technology
fjorgen@mit.edu

MS15

Long-Time Accuracy of Ensemble Kalman Filters with Machine-Learned Surrogate Models

The first part of this talk will study the long-time accuracy of ensemble Kalman filters. We introduce conditions on the dynamics and the observations under which the state estimation error remains small in the long-time horizon. Our theory covers a wide class of dissipative chaotic dynamical systems, which includes the Navier-Stokes equations and Lorenz models. In addition, we prove long-time accuracy of ensemble Kalman filters with surrogate dynamics, thus validating the use of machine-learned forecast models in ensemble data assimilation. The second part of this talk concerns the sample complexity of learning surrogate models for evolution equations with Fourier neural operators, motivated by the use of such surrogates in data assimilation tasks. Our results demonstrate that the sample complexity of learning a broad class of evolution equations is polynomial, not exponential, as is the case for generic Lipschitz operators. Further, our theory captures how the difficulty of learning depends on the structure of the PDE, including the degree of nonlinearity, the order of spatial derivatives, and the smoothness of solutions. This result helps to understand the success of operator learning techniques for learning PDEs and provides theoretical foundations for the use of neural-network surrogate models in filtering algorithms.

Nathan Waniorek
University of Chicago
waniorek@uchicago.edu

MS16

A Sequential Framework of Dimension-Reduced Expectation-Maximization with Optical Flow for Scalable Dynamic X-Ray Ct

Large-scale dynamic inverse problems are typically ill-posed and suffer from complexity of the model constraints and large dimensionality of the parameters. A common approach to overcome ill-posedness is through regularization that aims to add constraints on the desired parameters. In this talk, we describe a method for reconstructing a dynamic sequence of undersampled images from computerized tomography data. One approach is a prior-based dimension reduction Kalman filter with smoothing. However, while this method is suitable for image reconstruction, the challenge lies in the realization of the full model for applications, especially when the image sequence is non-stationary. To address this, we utilize a regularized Optical Flow and Expectation Maximization methods alongside the Kalman filter for robust construction of model realizations. We illustrate the effectiveness of the prescribed approach through numerical experiments from a wide range of limited angle and single shot computerized tomography.

Aryeh Keating
Virginia Polytechnic Institute and State University
aryehkeating@vt.edu

Mirjeta Pasha
Virginia Tech
mpasha@vt.edu

MS16

Prior-Driven Balancing for Bayesian Inverse Problems

Solving data assimilation (DA) problems typically requires expensive forward model simulations to compute the expected output for multiple initial conditions. Since this quickly becomes infeasible in high dimensions, reduced models can play a key role in rendering such computations tractable. In particular, model reduction methods from control (systems) theory are well established for obtaining efficient reduced models while preserving the map from control input to observed output. In this talk, we present prior-driven balancing as a system-theoretic model order reduction method adapted for use in DA. This includes a new interpretation of the Bayesian inverse problem, which is essential in DA, as a control problem and its reduction.

Josie König
University of Potsdam
josie.koenig@uni-potsdam.de

Elizabeth Qian
Georgia Institute of Technology
eqian@gatech.edu

Han Cheng Lie
University of Potsdam
han.lie@uni-potsdam.de

Melina A. Freitag
Universität Potsdam

melina.freitag@uni-potsdam.de

MS16

Likelihood-Informed Model Reduction for Bayesian Inference of Static Structural Loads

The Bayesian approach to solving inverse problems updates a prior to a posterior distribution of the parameters through data [Stuart and Dashti, *The Bayesian Approach to Inverse Problems*, 2017]. 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 inference 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 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 can be expressed as a low-rank update of the prior [Spantini et al, *Optimal Low-Rank Approximations of Bayesian Linear Inverse Problems*, 2015]. In this work, we introduce a framework that leverages the low-rank nature of the mapping between the data and the parameter itself to construct a reduced-order model for inverse problems. This approach is based on the eigendirections of the Fisher information and the prior precision and is referred to as likelihood-informed subspace (LIS). Through numerical examples, we compare the posterior approximation using the LIS reduced-order model with state-of-the-art approximation methods.

Jakob Scheffels
Technical University of Munich
jakob.scheffels@tum.de

Elizabeth Qian
Georgia Institute of Technology
eqian@gatech.edu

Iason Papaioannou
Engineering Risk Analysis Group
Technical University of Munich
iason.papaioannou@tum.de

Elisabeth Ullmann
Technical University of Munich
elisabeth.ullmann@tum.de

MS16

Goal-Oriented Dimension Reduction for Bayesian Inverse Problems Governed by Autonomous Dynamical Systems

We present a goal-oriented dimension reduction algorithm that extends our recent work on real-time, extreme-scale Bayesian inversion for linear autonomous dynamical systems. In that work, an offline-online decomposition enabled real-time predictions of time-dependent quantities of interest (QoIs) with uncertainties. However, the underlying inverse problem is extremely ill-posed given sparse and noisy data, and computing the full posterior covariance for the high-dimensional spatiotemporal parameter field remains computationally intractable. To overcome these challenges, we introduce a goal-oriented dimension reduction strategy that constructs a basis for the parameter field by retaining only the directions most informative to the QoIs. By leveraging the fast, FFT-based algorithms from our prior work for the parameter-to-observable and parameter-to-QoI maps, solving the full Bayesian inverse problem for this reduced parameter field becomes computationally viable. We establish theoretical error bounds for the proposed reduction strategy on model problems and demonstrate its effectiveness with numerical results for the acoustic-gravity model governing tsunami dynamics.

Sreeram R. Venkat
University of Texas Austin
201 E 24th St, Austin, TX 78712
srvenkat@utexas.edu

Stefan Henneking, Omar Ghattas
University of Texas at Austin
stefan@oden.utexas.edu, omar@oden.utexas.edu

MS17

Data-Driven Closures with Quantified Uncertainty to Enhance Model Prediction Reliability

Computational simulations of dynamical systems often require closure models to represent unresolved phenomena, enhance computational efficiency, or correct model-form error. While data-driven approaches offer a promising path to improving the accuracy and generality of closure models, a systematic, robust, and scalable methodology for constructing them has remained elusive. In this work, we introduce a novel framework: Scalable, Trustworthy, Agile, Goal-Oriented, Efficient, and Data-Assimilated Data-Driven Closure Models (STAGED DDCM). This two-step approach first decouples the forward computational model from closure learning via data assimilation, enabling flexible and scalable integration of observational data. In the second step, machine learning techniques are employed to learn the closure terms from the data generated via data assimilation, capturing complex, nonlinear relationships while preserving model efficiency and scalability. We demonstrate the end-to-end utility of the STAGED DDCM framework through two representative examples: (1) an epidemiological model, illustrating fast and accurate closure modeling for disease spread; and (2) a structural dynamics case involving a two-mass oscillator calibrated with real experimental data, showcasing the method's ability to capture unknown forces in predictive scenarios in physical systems where subject-matter expertise is limited. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Rileigh Bandy, Teresa Portone, Mohammad Khalil, Paul Smith, Eric Crislip, Kyle Neal
Sandia National Laboratories
rjbandy@sandia.gov, tporton@sandia.gov, mkhalil@sandia.gov, pausmit@sandia.gov, escrisl@sandia.gov, kneal@sandia.gov

MS17

A Score-Based Diffusion Model Approach for Adaptive Learning of Stochastic Partial Differential Equation Solutions

We propose a novel framework for adaptively learning the time-evolving solutions of stochastic partial differential equations (SPDEs) using score-based diffusion models within a recursive Bayesian inference setting. SPDEs play a central role in modeling complex physical systems under uncertainty, but their numerical solutions often suffer

from model errors and reduced accuracy due to incomplete physical knowledge and environmental variability. To address these challenges, we encode the governing physics into the score function of a diffusion model using simulation data and incorporate observational information via a likelihood-based correction in a reverse-time stochastic differential equation. This enables adaptive learning through iterative refinement of the solution as new data becomes available. To improve computational efficiency in high-dimensional settings, we introduce the ensemble score filter, a training-free approximation of the score function designed for real-time inference. Numerical experiments on benchmark SPDEs demonstrate the accuracy and robustness of the proposed method under sparse and noisy observations.

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

MS17

Conditional Diffusion Models for High-speed Boundary Layer Reconstructions

A complete resolution of turbulent boundary layers for high-speed flows is impossible for realistic applications in aerospace engineering. Direct numerical simulations (DNS) are available for canonical geometries and reveal information about the behavior of the near-wall flow. In this work, we use DNS data for high Mach turbulent Couette flow to train a diffusion-based super-resolution model for recovering the near-wall structure of the flow, given conditional information such as distance from the wall, Mach number, Reynolds number, etc. Our model demonstrates that the instantaneous flow structures of the boundary layer are recovered accurately, including derived quantities such as the wall-shear stress and wall heat flux with uncertainty estimates obtained through diffusion model sampling. This opens up a promising new pathway for the wall-modeling of large eddy simulations of turbulent boundary layers in hypersonic applications.

Hojin Kim
Purdue University
hojin.kim@purdue.edu

Dibyajyoti Chakraborty
Pennsylvania State University
d.chakraborty@psu.edu

Takahiko Toki, Carlo Scalo
Purdue University
ttoki@purdue.edu, cscalo@purdue.edu

Romit Maulik
Pennsylvania State University
rmm7011@psu.edu

MS17

Assessing the Quality of Uncertainty Estimates from Conformal Classification Methods

Conformal prediction is becoming widely used as a distribution-free approach to quantifying uncertainty. The minimal assumptions, flexibility in application to different types of models and data, and guarantee of valid uncertainty estimates make conformal methods attractive for researchers and practitioners alike. Because, under some weak assumptions, conformal methods are guaranteed to produce valid uncertainty quantification (UQ), the quality of the resulting UQ is typically assessed by simply considering efficiency. For two conformal approaches that both produce valid UQ, the smaller interval or set is more efficient. While a straightforward concept, the assessment of efficiency presents some challenges. In this talk, we present an overview of conformal prediction for classification and discuss how to compare conformal methods in terms of UQ quality. We introduce several novel measures of efficiency and compare them to existing approaches through simulation studies. We argue that this represents a more holistic approach to efficiency assessment than current common practices. Finally, we discuss how our approach to evaluating efficiency using simulation studies may be carried out with real data sets and provide some potential avenues for future work in this direction.

Jason Adams, Shengzhu Yin, Brandon Berman
Sandia National Laboratories
jradams@sandia.gov, syin@sandia.gov,
bjberma@sandia.gov

MS18

Generalized Least Volume for Structured Representation Learning

We introduce Generalized Least Volume (GLV), a regularization framework for autoencoders that extends geometric ideas of latent volume minimization beyond Euclidean settings. GLV incorporates label or side information into the latent representation, yielding contrastive effects for discrete labels and smooth interpolations for continuous ones. We show that GLV preserves the dimension-reduction benefits of Least Volume while providing theoretical guarantees under Lipschitz continuity assumptions. Empirically, GLV reveals dataset-specific topological structure and enhances stability in downstream tasks: on labeled benchmarks it induces label-aware disentanglement, while on a continuous-label airfoil dataset it produces representations that correspond to smooth aerodynamic variations, enabling more robust optimization. These results highlight GLV as a principled approach to integrating geometric and statistical structure in latent-variable models.

Qiuyi Chen
University of Maryland
Qiuyi.Chen@gevernova.com

Mark Fuge
ETH Zurich
mafuge@ethz.ch

Cashen Diniz
University of Maryland
cdiniz@umd.edu

MS18

Feature-to-feature Data Assimilation: Application to High-speed Reactive Flow

We describe recent algorithms to evaluate the distribution of spatio-temporal features conditional on samples of other, related, features. In particular, we evaluate the flow morphology inside a channel with high-speed reactive flow given temporal observations at sparse boundary locations. The approach consists of a reverse diffusion process conditioned by a polynomial chaos expansion. The PCE encodes the joint dependence between features and signals and the

reverse diffusion process serves to filter and sharpen the features.

Zheming Gou
University of Southern California
zgou@usc.edu

Roger Ghanem
University of Southern California
Civil and Environmental Engineering
ghanem@usc.edu

MS18

From Smooth to Nonsmooth: Minimax Optimal Regression with Laplacian Eigenmaps

In this talk, I will discuss nonparametric regression under varying smoothness assumptions, ranging from classical Sobolev spaces to fractional, nonsmooth regimes. In the smooth setting, we develop Laplacian Eigenmapbased estimators and establish both adaptive and non-adaptive minimax optimality. Our analysis covers a broad family of weighted graph Laplacians and shows that principal components regression on their eigenfunctions achieves the optimal rate, with adaptation accomplished via a Lepskitype procedure. In the nonsmooth setting, where the regression function belongs to a fractional Sobolev space, we introduce Principal Component Regression with Fractional Laplacian Eigenmaps (PCRFLE). We prove that this estimator attains the minimax optimal rate for fractional smoothness, and complement the upper bounds with matching lower bounds. Taken together, these results provide a unified spectral framework for minimax optimal regression via Laplacian embedding, bridging the gap between smooth and nonsmooth function classes and providing a theoretical foundation for (fractional) Laplacian involved high-dimensional nonparametric learning.

Zhaoyang Shi
Harvard University
zshi@fas.harvard.edu

Krishnakumar Balasubramanian
University of California, Davis, U.S.
kbala@ucdavis.edu

Wolfgang Polonik
University of California Davis
wpolonik@ucdavis.edu

MS18

Unsupervised Neural Operators for Optimal Control: Scaling Laws and Applications

Optimal control provides a principled framework for transforming dynamical system models into intelligent decision-making, yet classical computational approaches are often too expensive for real-time deployment in dynamic or uncertain environments. In this work, we propose a neural operator learning framework designed for optimal control problems. It offers a new paradigm by directly approximating the mapping from system conditions to optimal control strategies, enabling instantaneous inference across diverse scenarios once trained. We investigate both the promise and limitations of neural operators for optimal control. Theoretically, we establish scaling laws that quantify how sample and model complexity depend on the intrinsic dimension and regularity of the optimal control function, thus clarifying the price to pay for efficiency. Numerically, we demonstrate case studies where neural operators achieve remarkable gains in real-time performance, as well as regimes where complexity barriers persist. Together, these results provide a balanced perspective: neural operators represent a powerful new tool for high-performance control when hidden low-dimensional structure can be exploited, yet remain fundamentally constrained by dimensional complexity in more challenging settings.

Wuzhe Xu, Rongjie Lai
Purdue University
xu2224@purdue.edu, lairj@purdue.edu

MS19

Particle Rates for Stein Variational Gradient Descent

Stein Variational Gradient Descent (SVGD) is an algorithm for sampling from a target distribution using a deterministic flow of N particles (started from a possibly random initial distribution) which interact with each other through a positive-definite kernel. In comparison with traditional Markov Chain Monte Carlo schemes, SVGD requires less randomness, is easier to implement and is seen to empirically perform better in many situations. SVGD can be thought of as a 'discretization of a kernelized gradient flow for the relative entropy on Wasserstein-2 space. Thus, its convergence properties are well understood in the mean-field limit. However, due to relative entropy having no suitable analogue for the empirical measure of particle locations, obtaining rates of convergence for the actual particle system has been a long-standing open problem. The few previously obtained rates of order $1/\sqrt{\log(\log(N))}$ were believed to be highly sub-optimal. In the current work, we make the observation that the relative entropy of the joint distribution of particle locations with respect to the N-fold product of the target distribution can be used as a 'Lyapunov function which furnishes near-optimal rates of convergence of order $1/\sqrt{N}$. We also obtain long-time propagation of chaos results and marginal convergence rates for the particle laws.

Sayan Banerjee
University of North Carolina, Chapel Hill
sayan@email.unc.edu

MS19

Using Probability Paths to Improve the Efficiency of Sampling and Estimation Problems

Modern algorithms for sampling, estimating normalizing constants, and density-ratio estimation often rely on a probability path that connects a complex target distribution to a simple base distribution such as a Gaussian. In this talk, I will discuss the limitations of classical methods, show how using probability paths can improve their efficiency, sometimes dramatically, and present principled strategies for designing the probability path.

Omar Chehab
Carnegie Mellon University
ochehab@andrew.cmu.edu

MS19

Uniform-in-Time Propagation of Chaos for Consensus-Based Optimization

Consensus-Based Optimization (CBO) is a class of

derivative-free global optimization algorithms governed by interacting stochastic particle systems. These systems are characterized by a non-local structure where particles are attracted toward a weighted consensus point via a Laplace-principle-based approximation of the global minimizer. This mechanism allows for global optimization without gradient information, and causes mathematical challenges for a rigorous mean-field analysis, necessitating new approaches. ¿ ¿ In this talk, we establish uniform-in-time propagation of chaos for CBO. By leveraging a novel stability estimate for the weighted mean and quantitative concentration inequalities, we prove that the microscopic particle system tracks the mean-field limit with the classical Monte Carlo rate over infinite time horizons. Notably, for CBO with anisotropic noise, the error prefactor is independent of the problem dimension. We also present results on the almost sure exponential convergence of the microscopic system.

Dohyeon Kim
California Institute of Technology
dohyeon@caltech.edu

Nicolai Gerber
Ulm University, Germany
nicolai.gerber@uni-bonn.de

Franca Hoffmann
California Institute of Technology
franca.hoffmann@caltech.edu

Urbain Vaes
École des Ponts and Inria Paris
urbain.vaes@inria.fr

MS19

Blind Denoising Diffusion Models and the Blessings of Dimensionality

We analyze, theoretically and empirically, the performance of generative diffusion models based on *blind denoisers*, in which the denoiser is not given the noise amplitude in either the training or sampling processes. Assuming that the data distribution has low intrinsic dimensionality, we prove that blind denoising diffusion models (BDDMs), despite not having access to the noise amplitude, *automatically* track a particular *implicit* noise schedule along the reverse process. Our analysis shows that BDDMs can accurately sample from the data distribution in polynomially many steps as a function of the intrinsic dimension. Empirical results corroborate these mathematical findings on both synthetic and image data, demonstrating that the noise variance is accurately estimated from the noisy image. Remarkably, we observe that schedule-free BDDMs produce samples of higher quality compared to their non-blind counterparts. We provide evidence that this performance gain arises because BDDMs correct the mismatch between the true residual noise (of the image) and the noise assumed by the schedule used in non-blind diffusion models. This is joint work with Zahra Kadkhodaie, Sinho Chewi, and Eero Simoncelli.

Aram-Alexandre Pooladian
New York University
aram-alexandre.pooladian@yale.edu

MS20

Flow Map Learning of Unknown Stochastic Systems

We present a computational framework for modeling unknown stochastic systems using observation data. The method utilizes data on the system trajectories and constructs a stochastic flow map of the unknown system. The stochastic flow map takes the form of generative model, e.g., GANs, autoencoder, normalizing flow, etc. Once the generative stochastic flow map is constructed, it allows us to create a predictive model of the known stochastic system. A set of numerical examples, including complex chemical reaction networks, are presented to demonstrate the effectiveness of the method. By using short bursts of data, the method is able to construct accurate models for the unknown system for very long-term system prediction.

Yuan Chen
Ohio State University
chen.11050@osu.edu

MS20

To Be Announced

TBD

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@eoas.ubc.ca

MS20

Goal-Oriented Inference with Conditional Generative Models

High-fidelity numerical simulation is essential for predicting unobservable quantities of interest (QoI) in physical phenomena. However, solving goal-oriented inverse Bayesian inference problems that condition predictions on data remains computationally challenging, as they often require sampling from a high-dimensional posterior and simulating the prediction model at these samples to compute the QoI. In this talk, we present conditional generative modeling techniques for efficient goal-oriented inference that directly condition QoI predictions on data, bypassing explicit posterior sampling. Through extensive benchmarking with existing methods, we demonstrate that our goal-oriented generative models can leverage the inherent low-dimensional structure of the parameter-to-QoI map to construct computationally efficient solutions to inverse problems with large numbers of uncertain variables.

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

MS21

Map: Multi-Human-Value Alignment Palette for Safe and Personalized Generative AI

As AI systems transition from automation tools to autonomous collaborators, building user trust remains a fundamental challenge: users face repeated misalignment with their preferences, unreliable outputs mixed with hallucinations, and sycophantic responses lacking critical thinking. This talk presents an approach for enhancing trustworthy human-AI coworking systems.

Jie Ding
University of Minnesota Twin Cities

dingj@umn.edu

MS21

GaussNet: A Gaussian Process Network Accelerated by the Inverse Kalman Filter

In the first part of the talk, we introduce the inverse Kalman filter, which enables exact matrix-vector multiplication between a covariance matrix from a dynamic linear model, including the Matrn covariance matrix with a half-integer roughness parameter, and a real-valued vector with a linear computational cost to the dimension of the vector. We integrate the inverse Kalman filter with the conjugate gradient algorithm, which substantially accelerates the computation of matrix inversion for an additive summation of covariance matrices, where other approximation approaches may not be directly applicable. In the second part of the talk, we introduce GaussNet, a new neural network architecture that requires much less data for predicting nonlinear maps than conventional neural networks. We derive the closed-form expressions of the predictive distribution and gradients to train GaussNet, which can be substantially accelerated by the inverse Kalman filter. We demonstrate the scalability and accuracy of GaussNet for predicting nonlinear relationships with general multi-dimensional inputs compared with approximation approaches of Gaussian processes and neural networks. Other applications include nonparametric estimation of particle interaction functions using trajectories of cells from simulation and microscopy videos.

Mengyang Gu
Department of Statistics and Applied Probability
University of California, Santa Barbara
mengyang@pstat.ucsb.edu

MS21

Embedding Machine Learning into a Numerical Cimate Models: Computational, Probabilistic and Scientific Advances

Embedding Machine Learning into a numerical climate models: computational, probabilistic and scientific advances This talk addresses the design, development and deployment of using Gaussian processes to improve convection within a climate model. The climate model's temperature and humidity profiles are stochastically perturbed every 6 hours at runtime based on a Gaussian process trained on data from a high-resolution weather model with realistic convection. A proof-of-concept study shows the effects of the approach on a simplified climate model's output (SPEEDY). The precipitation pattern is improved globally with the most significant reductions in the Indian Ocean and tropical Pacific Ocean off the Mexican coast which yields a reduced area-weighted RMSE in the tropics by 20% (Giles et al. 2024). A second project looks at the engineering challenges of compiling the Gaussian process deployment into a Fortran-based climate model at runtime with minimal overhead. Initial findings conclude that running on an NVIDIA GH200 device reduces the overheads from 3x to only 3% in compute time but yields similar modelling improvements. This extension also uses a more sophisticated climate model, CESM 2.1, with a moderate resolution. Finally we will discuss the first steps towards including coastal wave modelling into a climate model, using Smoothed Particle Hydrodynamics as well as air-sea exchanges representation.

James Briant
University College London
james@briant.co.uk

debiprasad bhuyan, Saroj Mishra
IIT Delhi
debiprasadbhuyaniit@gmail.com, skm@cas.iitd.ac.in

Shipra Jain, Serge Guillas
University College London
shipra.jain@ucl.ac.uk, s.guillas@ucl.ac.uk

MS21

FunDiff: Diffusion Models over Function Spaces for Physics-Informed Generative Modeling

Recent advances in generative modeling – particularly diffusion models and flow matching – have achieved remarkable success in synthesizing discrete data such as images and videos. However, adapting these models to physical applications remains challenging, as the quantities of interest are continuous functions governed by complex physical laws. Here, we introduce FunDiff, a novel framework for generative modeling in function spaces. FunDiff combines a latent diffusion process with a function autoencoder architecture to handle input functions with varying discretizations, generate continuous functions evaluable at arbitrary locations, and seamlessly incorporate physical priors. These priors are enforced through architectural constraints or physics-informed loss functions, ensuring that generated samples satisfy fundamental physical laws. We theoretically establish minimax optimality guarantees for density estimation in function spaces, showing that diffusion-based estimators achieve optimal convergence rates under suitable regularity conditions. We demonstrate the practical effectiveness of FunDiff across diverse applications in fluid dynamics and solid mechanics. Empirical results show that our method generates physically consistent samples with high fidelity to the target distribution and exhibits robustness to noisy and low-resolution data.

Sifan Wang, Lu Lu
Yale University
sifan.wang@yale.edu, lu.lu@yale.edu

MS22

Hybrid Data Assimilation and Deep Learning Approaches for Sparse Observational Data: Applications in Air Pollution and Geosciences

Reconstructing spatiotemporal systems from sparse data remains a major challenge in geoscience, air pollution, and fluid dynamics. Data assimilation (DA) methods are often computationally expensive for multivariate, high-dimensional systems. Machine learning (ML) models struggle with sparse and shifting sensor placements. The reconstruction problem is also ill-posed, which increases the risk of overfitting. We present recent work addressing these challenges through latent DA algorithms that reduce the computational cost of variational DA by assimilating sparse observations in a neural networkbased reduced-order latent space, enabled by the GPU-accelerated TorchDA library. In air pollution applications, this framework supports co-assimilation of NO2, O3, and PM fields within a shared latent space, improving multivariate consistency. We further employ advanced deep learning approaches, including Voronoi-tessellation CNNs and Vision Transformer autoencoders, to reconstruct full spatial fields from sparse, irregular, and movable sensors, with integration into DA

frameworks also evaluated. Finally, we explore generative AI methods, particularly denoising diffusion models, which sample from probability distributions to mitigate overfitting and act as ensemble predictors for uncertainty quantification. Results span fluid dynamics benchmarks and semi-operational air pollution simulations.

Sibo Cheng
CEREA, Ecole des ponts, Institut Polytechnique de Paris
sibo.cheng@enpc.fr

Sabnis Abhishek, Mihai Mitrea
Imperial College London
abhishek.sabnis24@imperial.ac.uk,
mihai.mitrea24@imperial.ac.uk

Lya Lugon
ENPC
lya.lugon@enpc.fr

Karine KATA-SARTELET
CEREA, Ecole des ponts, Institut Polytechnique de Paris
karine.sartelet@enpc.fr

Marc Bocquet
Ecole des Ponts ParisTech
marc.bocquet@enpc.fr

MS22

Can the Generative AI Shark Eat the Entire Data Assimilation System?

The recent surge in machine learning (ML) methods for geophysical modeling has raised the question of how these methods might be applied to data assimilation (DA). We focus on diffusion modeling (a form of generative artificial intelligence) for systems that can perform the entire DA process, rather than on ML-based tools used within a conventional DA system. We identify at least three distinct types of diffusion-based DA systems and show that they differ in the posterior distribution they target for sampling. These posterior distributions correspond to different priors and/or likelihoods, which in turn result in unique training datasets, computational requirements, and state estimate qualities. Our analysis further shows that a diffusion DA system designed to target the same posterior distribution as current ensemble DA algorithms requires re-training at each DA cycle, which is computationally costly. We discuss the implications of these findings for the use of diffusion modeling in DA.

Matthias Morzfeld
University of California, San Diego
Scripps Institution of Oceanography
mmorzfeld@ucsd.edu

Daniel Hodyss
U.S. Naval Research Laboratory
daniel.hodyss@nrl.navy.mil

MS22

AI-Boosted Rare Event Sampling to Reduce Extreme Weather Uncertainty

Estimating the risk of extreme weather due to climate change, especially at regional scales, is critically important yet remains one of the most challenging questions to investigate with traditional physics-based global climate models (GCMs). The computational cost of these models leads to high uncertainty in estimates of the rarest (and most impactful) extreme-weather return times. AI emulators trained on historical reanalysis to emulate global-atmosphere dynamics have demonstrated high forecast accuracy, greatly reduced computational cost, and decades-long stability. A natural question, therefore, is whether these models can reduce the uncertainty of very rare weather events. Because this is difficult to assess with only the historical record, we simulate 92,000 years of a stationary climate, train two stable AI emulators using 100 years of data, and generate 92,000-year emulation ensembles from each emulator. While we find that emulators can generate extreme events beyond their training data and similar to the most extreme events in the long simulation, they do not exactly reproduce the return period of extreme heat, cold, and precipitation in every region. To correct this, we introduce AI+RES. In this method, we use the emulated ensemble forecast as the score function in a rare-event sampling algorithm. Using AI+RES, we can predict unbiased extreme heat spell return times an order of magnitude less uncertainty than direct simulation with the true GCM.

Alexander Wikner
University of Chicago
awikner1@umd.edu

Amaury Lancelin
École Normale Supérieure
Laboratoire de Météorologie Dynamique
lancelin.amaury@lmd.ipsl.fr

Troy Arcomano
AI2
troya@allenai.org

Karan Jakhar
University of Chicago
karanj@uchicago.edu

Dhruvit Patel
Silurian AI
dhruvit@silurian.ai

Dorian S. Abbot
Department of Geophsyical Sciences
University of Chicago
abbot@uchicago.edu

Freddy Bouchet
Laboratoire de Météorologie Dynamique
CNRS
freddy.bouchet@lmd.ipsl.fr ¡freddy.bouchet@lmd.ips

Pedram Hassanzadeh
The University of Chicago
pedramh@uchicago.edu

Jonathan Weare
Courant Institute of Mathematical Sciences
weare@cims.nyu.edu

MS23

Efficient Importance Sampling for Value of Information Analysis in the Presence of Rare Events

In decision-making situations characterized by uncertainty

and risk, there often exist opportunities to gather additional information, contributing to better informed decisions. The so-called Value of Information (VoI) can be evaluated to enhance data acquisition, but its major drawback lies in its high computational complexity. VoI evaluation necessitates integration over potential observation outcomes for each candidate strategy, incorporating optimization over possible actions based on the expected costs for the action-measurement combinations. We focus on efficient estimation of the outermost integration over potential observation outcomes. Estimation of this integral with brute-force Monte Carlo sampling is inefficient in cases with several measurements having more than two possible outcomes. We propose a novel importance sampling scheme, which sequentially approximates the optimal importance sampling density using a Gaussian process model of the logarithmically-transformed failure probability. Thereby, the sequentially drawn training points and their corresponding failure probability estimates serve as both importance samples and updates to the surrogate model, with accurate estimation of low-probability outcomes being essential. Using engineering test cases, we demonstrate the methods potential to accelerate VoI calculations, with the hope that their implementation can enable more frequent usage in practice.

Lea Friedli
University of Bern
lea.friedli@hotmail.com

Iason Papaioannou
Engineering Risk Analysis Group
Technical University of Munich
iason.papaioannou@tum.de

Daniel Straub
Engineering Risk Analysis Group
TU München
straub@tum.de

MS23

Advances in Importance Sampling for Rare Event Probability Estimation

Rare event simulation and probability estimation for safety-critical systems under computationally expensive simulators is a central UQ task. Several variance-reduction methods, based on subset/splitting, cross-entropy, sequential Monte Carlo, etc., can mitigate the computational cost. However, in high-dimensional spaces with strongly nonlinear forward maps, disconnected modes, and general input distributions, available methods often fail to perform adequately and efficiently. We present here a novel two-stage importance sampling (IS) estimator addressing these challenges. A flexible, unnormalized target is first constructed, in a general fashion for all cases. It intrinsically concentrates probability mass on the domain of interest, while relaxing the optimal yet intractable IS density, enabling effective sampling with diverse MCMC methods in gradient-based and gradient-free settings, e.g., QNp-HMCMC and pCN. The normalizing constant of the approximate target is then estimated using density estimators trained on the MCMC samples. Neural density estimators can further boost performance, particularly our Score Aligned Neural Density Estimator (SANDE). SANDE aligns the score field of the learned density with that of the unnormalized target, enabling accurate and scalable density estimation at competitive cost. The proposed framework is validated on challenging benchmarks and applications of complex autonomous and engineered systems involving several hundred random variables.

Kostas G. Papakonstantinou
The Pennsylvania State University
kpapakon@psu.edu

Elsayed Eshra
Pennsylvania State University
eme5375@psu.edu

MS24

Worst-Case Learning of Adverse Quantities of Interest

After a brief rundown on Optimal Recovery, viewed as an analytical learning theory focused on worst-case guarantees, I will present new works whose common theme is the estimation of quantities of interest that are not simple linear maps. In one direction, the quantity of interest is a nonlinear functional (e.g. the maximum of a function, a piecewise affine functional, etc.), but recovery procedures that are genuinely worst-case optimal can be constructed. In another direction, the quantity of interest may be linear but is not explicit (e.g. the solution operator in partial differential equations), so the existing theory does not directly apply.

Simon Foucart
Texas A&M University
foucart@tamu.edu

MS24

Structure-Preserving Discontinuous Galerkin Methods for Stochastic Shallow Water Equations

Shallow water equations (SWE) are fundamental models in fluid dynamics that are essential for studying a wide range of geophysical and engineering phenomena. Therefore, stable and accurate numerical methods for SWE are needed. Although some algorithms are well studied for deterministic SWE, more effort should be devoted to handling the SWE with uncertainty, e.g., from initial value and bottom topography. Building from the hyperbolicity-preserving stochastic Galerkin formulation for SWE [Dai, Epshteyn, Narayan, 2021 SISC] and the entropy stable discontinuous Galerkin for skew-symmetric SWE [Fu, 2022 JSC], we develop a structure-preserving, entropy conservative, and entropy stable discontinuous Galerkin method for the stochastic Galerkin shallow water system. We demonstrate the efficacy, applicability and robustness of these structure-preserving algorithms through several numerical experiments. Keywords: stochastic shallow water equation, stochastic Galerkin, discontinuous Galerkin, structure-preserving, entropy stable.

Yinqian Yu
University of Utah
yinqian.yu@utah.edu

MS25

Hard-Constrained Neural Networks for Learning Chaotic PDEs with Boundedness Guarantees

Machine learning models often fail to respect known structure, such as physical laws, and soft-constraining penalty methods are often inadequate, especially in out-of-distribution scenarios. We present a new paradigm: hard-

constrained neural networks (HardNet). HardNet appends a differentiable enforcement layer to a networks outputs to guarantee satisfaction of input-output constraints during both training and inference. It preserves the expressive power of neural networks and allows for end-to-end training with standard unconstrained optimization methods. I will then specialize the idea to learning chaotic partial differential equations, where existing machine learning models often generate unbounded trajectories, leading to invalid statistics. We jointly learn a dissipative dynamics emulator and an energy-like function whose algebraic conditions imply convergence to an invariant level set. Enforcing those conditions via an explicit convex quadratic projection layera HardNet-style enforcement layerprovides formal boundedness guarantees and an outer estimate of the strange attractor. On KuramotoSivashinsky and Navier-Stokes benchmarks, the method delivers bounded long-horizon forecasts and, as a result, preserves invariant statistics. Together, these results show the effectiveness of enforcement layers as a general recipe for structure-preserving learning with provable guarantees and strong empirical performance in scientific, or otherwise high-stakes, applications.

Navid Azizan
MIT
azizan@mit.edu

Andrea Goertzen
Massachusetts Institute of Technology
agoertz@mit.edu

MS25

Semi-Explicit Neural DAEs: Learning Long-Horizon Dynamical Systems with Algebraic Constraints

Despite the promise of scientific machine learning (SciML) in combining data-driven techniques with mechanistic modeling, existing approaches for incorporating hard constraints in neural differential equations (NDEs) face significant limitations. Scalability issues and poor numerical properties prevent these neural models from being used for modeling physical systems with complicated conservation laws. We propose Manifold-Projected Neural ODEs (PNODEs), a method that explicitly enforces algebraic constraints by projecting each ODE step onto the constraint manifold. This framework arises naturally from semi-explicit differential-algebraic equations (DAEs), and includes both a robust iterative variant and a fast approximation requiring a single Jacobian factorization. We further demonstrate that prior works on relaxation methods are special cases of our approach. PNODEs consistently outperform baselines across six benchmark problems achieving a mean constraint violation error below 10^{-10}. Additionally, PNODEs consistently achieve lower runtime compared to other methods for a given level of error tolerance. These results show that constraint projection offers a simple strategy for learning physically consistent long-horizon dynamics.

Avik Pal
Massachusetts Institute of Technology
avikpal@mit.edu

MS25

Quantifying Model Uncertainties in Exact Conservation Law-Embedded Reduced Dynamics Models

Reduced/Latent state dynamics approaches for parameterized PDEs offer a viable method where high-fidelity data and machine learning techniques are used to reduce the system's dimensionality and estimate the dynamics of low-dimensional reduced states. These reduced state dynamics approaches rely on high-quality data and struggle with highly sparse spatiotemporal noisy measurements typically obtained from experiments. Furthermore, there is no guarantee that these models satisfy governing physical conservation laws, especially for parameters that are not a part of the model learning process. We propose a reduced state dynamics approach, which we refer to as ECLEIRS, that satisfies conservation laws exactly even for parameters unseen in the model training process. We compare ECLEIRS with other reduced state dynamics approaches, those that do not enforce any physical constraints and those with physics-informed loss functions, for three shock-propagation problems: 1-D advection, 1-D Burgers and 2-D Euler equations. The numerical experiments conducted in this study demonstrate that ECLEIRS provides the most accurate prediction of dynamics for unseen parameters even in the presence of highly sparse and noisy data. We also demonstrate that ECLEIRS yields solutions and fluxes that satisfy the governing conservation law up to machine precision for unseen parameters, while the other methods yield much higher errors and do not satisfy conservation laws.

Aviral Prakash
Los Alamos National Laboratory
aviralp@lanl.gov

MS26

Optimal Tensor Algebras Via Variable Projection

As high-dimensional, complex data becomes more accessible, tensor-based approaches become increasingly attractive for multiway data analysis. However, many multilinear methods suffer from a so-called curse of multidimensionality; that is, that fundamental mathematical properties break down in high dimensions. Recent advances in matrix-mimetic tensor algebra in have made it possible to preserve linear algebraic properties and, as a result, to obtain optimal representations of multiway data. Matrix mimeticity arises from interpreting tensors as operators, which in turn are parameterized by invertible linear transformations. The choice of transformation is critical to representation quality and computational efficiency, and thus far, has been made heuristically. In this talk, we will learn data-dependent, orthogonal transformations by leveraging the optimality of matrix-mimetic representations. In particular, we will exploit the coupling between transformations and optimal tensor representations using variable projection. We will highlight the efficacy of our proposed approach on multiway data compression and regression tasks.

Elizabeth Newman
Emory University
e.newman@tufts.edu

MS26

The Tensor-Train Stochastic Finite Volume Method for Uncertainty Quantification

Many problems in physics and engineering are modeled by systems of partial differential equations such as the shallow water equations of hydrology, the Euler equations for invis-

cid, compressible flow, and the MHD equations. The initial data, boundary conditions, and coefficients of these models may be uncertain due to measurement, prediction, or modeling errors. The stochastic finite volume (SFV) method offers an efficient one-pass approach for assessing uncertainty in hyperbolic conservation laws. The SFV method has shown great promise as a weakly-intrusive PDE solver for uncertainty quantification. However, in many relevant applications, the dimension of the stochastic space can make traditional implementations of the SFV method infeasible or impossible due to the so-called curse of dimensionality. We introduce the Tensor-Train SFV (TT-SFV) method within the tensor-train framework to manage the curse of dimensionality. This integration, however, comes with its own set of difficulties, mainly due to the propensity for shock formation in hyperbolic systems. To overcome these issues, we have developed a tensor-train-adapted stochastic finite volume method that employs a global WENO reconstruction, making it suitable for such complex systems. This approach represents the first step in designing efficient tensor-train techniques for uncertainty quantification in hyperbolic systems and conservation laws involving shocks.

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

Steven Walton
Los Alamos National Laboratory
stevenw@lanl.gov

Gianmarco Manzini
Los Alamos National Laboratory, IMATI-CNR
marco.manzini@imati.cnr.it

MS27

Sensitivities, Singular Values, and Optimally Learned Quantities of Interest

Traditional approaches for optimal experimental design (OED) typically require solving a large number of computationally intensive inverse problems to find the data that maximizes a utility function. We introduce two novel OED criteria specifically crafted for the data consistent inversion (DCI) framework that avoid approximating inverse problem solutions. DCI is a specific approach for solving a class of stochastic inverse problems by constructing a pullback measure on uncertain parameters from an observed probability measure on the outputs of a quantities of interest (QoI) map. While expected information gain (EIG) has been used for both DCI and Bayesian based OED, the characteristics and properties of DCI solutions differ from those of solutions to Bayesian inverse problems which should be reflected in the OED criteria. The new design criteria developed in this study, called the expected scaling effect and the expected skewness effect, leverage the geometric structure of pre-images associated with observable data sets, allowing for an intuitive and direct approach to OED. These criteria utilize singular value computations derived from sampled and approximated Jacobians of the experimental designs. We present both simultaneous and sequential (greedy) formulations of OED based on these innovative criteria. Numerical results demonstrate the effectiveness in our approach for DCI solutions to stochastic inverse problems.

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS27

Bayesian D-Optimal Experimental Designs Via Column Subset Selection

This work tackles optimal sensor placement for Bayesian linear inverse problems, a popular version of the more general Optimal Experimental Design (OED) problem, using the D-optimality criterion. This is done by establishing connections between sensor placement and Column Subset Selection Problem (CSSP), which is a well-studied problem in Numerical Linear Algebra (NLA). In particular, we use the Golub-Klema-Stewart (GKS) approach which involves computing the truncated Singular Value Decomposition (SVD) followed by a pivoted QR factorization on the right singular vectors. The algorithms are further accelerated by using randomization to compute the low-rank approximation as well as for sampling the indices. The resulting algorithms are robust, computationally efficient, amenable to parallelization, require virtually no parameter tuning, and come with strong theoretical guarantees. One of the proposed algorithms is also adjoint-free which is beneficial in situations, where the adjoint is expensive to evaluate or is not available. Additionally, we develop a method for data completion without solving the inverse problem. Numerical experiments on model inverse problems involving the heat equation and seismic tomography in two spatial dimensions demonstrate the performance of our approaches.

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Srinivas Eswar
Mathematics and Computer Science Division
Argonne National Laboratory
seswar@anl.gov

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS28

Estimating Lung Stiffness from CT Using PDE-Constrained Learning

Tissue elasticity tells us a great deal about how diseases begin and progress, but measuring it inside the body isnt always straightforward. Techniques like Magnetic Resonance Elastography have made it possible to non-invasively assess tissue stiffness in organs such as the liver, yet the same approach struggles in the lungs because of their complex anatomy and air-filled structure. In this talk, Ill introduce our effort to estimate lung stiffness directly from standard CT scans acquired during regular breathing. The idea is to link image-derived lung motion to underlying tissue mechanics. We start by estimating how the lung deforms between respiratory phases using image registration, then use a physics-based model to relate that deformation to tissue elasticity. Instead of solving this challenging inverse problem separately for each patient, we train a deep learning model to learn this mapping. Importantly, the model is guided by physics: its predictions must reproduce observed deformations when plugged into a physical simulation. Ill discuss how this physics-constrained learning framework can recover spatial patterns of lung stiffness and what it reveals about the feasibility of performing virtual

elastography directly from CT.

Kayhan Batmanghelich
Boston University
batman@bu.edu

MS28

Amortized Inference of Physiologic Boundary Conditions

Amortized probabilistic inference uses neural networks to learn parametric families of posterior distributions. Despite an increased offline training cost, online generation of posterior samples is extremely fast and, unlike traditional Markov Chain Monte Carlo (MCMC) methods, the system does not need to be retrained every time new observations become available. Amortized inference is therefore particularly interesting for patient-specific cardiovascular modeling, where boundary condition tuning is a necessary prerequisite that can quickly become intractable when combining three-dimensional cardiovascular models with sample-based approaches to inverse problems. Given high clinical demands and limited computational resources typically available, it is common practice to use low-fidelity (zero- or one-dimensional) models to capture flow and pressure results along a centerline or at discrete locations in lumped-parameter circuit models. We introduce a simulation-free amortized inference framework based on conditional density estimation. The framework predicts the posterior distribution of patient-specific boundary conditions using neural flow models conditioned on embeddings of complex patient-specific anatomies, in addition to typical clinical targets. We test the framework on patient-specific models that are readily available through the open source Vascular Model Repository (VMR). *Supported by Yansouni Family Stanford Graduate Fellowship and NSF grant 2105345.

Chloe Choi
Stanford University
chloe1@stanford.edu

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

Alison L. Marsden
Stanford University
amarsden@stanford.edu

MS28

Surrogate Models for Handling Global and Local Shape Variations in Aortic Hemodynamics

This work addresses the efficient handling of shape variability in aortic geometries, where we distinguish between local and global shape variations. Local shape variations can be used to describe uncertainty from medical imaging and the algorithm chain that leads to a simulation-ready discrete representation. Global shape variations capture the inter-patient variability (see, e.g., the SynthAorta dataset [Bonjak et al., SynthAorta: A 3D Mesh Dataset of Parametrized Physiological Healthy Aortas, IEEE Trans. Med. Imag., 2025]). Our investigations are based on modeling the geometrical perturbation and the resulting mesh generation, combined with computational fluid dynamics simulation. Variability across shapes is handled using an automated registration based on ResNet-LDDMM. This approach allows to use the simulation data to train shape-informed graph neural network models [Romor et al., Data assimilation performed with robust shape registration and graph neural networks: application to aortic coarctation, arxiv, 2025] to create efficient surrogates. The approach can be used to quantify the effect of shape variability on relevant biomarkers associated with clinical outcomes in numerous vascular pathologies, reducing the computational costs associated with numerical simulations.

Richard Schussnig
Graz University of Technology
Austria
richard.schussnig@gmx.at

Francesco Romor
SISSA, International School for Advanced Studies
Trieste, Italy
francesco.romor@wias-berlin.de

Sascha Ranftl
Graz University of Technology
Institute of Theoretical Physics-Computational Physics
sr7843@nyu.edu

Alfonso Caiazzo
Weierstrass Institute for Applied Analysis and Stochastics
Berlin
caiazzo@wias-berlin.de

Gerhard A. Holzapfel
Institute of Biomechanics, Graz University of Technology
holzapfel@tugraz.at

MS28

Reconstructing Boundary Waveform of Cerebrospinal Fluid Flow with Coupled Fluid-Solid Simulation and Surrogate Modeling

Intrathecal injection delivers medication directly to the central nervous system, with transport governed by cerebrospinal fluid (CSF) flow within the thecal space. CSF flow is periodic and influenced by the stiffness of surrounding soft tissues, with waveform shape and period exhibiting significant patient variability that can affect drug delivery. Although coupled fluidsolid solvers in spinal geometries can accurately model CSF dynamics, their high computational cost limits their use in inverse analysis. Here, we develop a surrogate model of CSF flow that maps the foramen magnum (FM) inlet boundary condition and elastic tissue properties directly to displacement and velocity fields. FM boundary conditions are sampled using a Gaussian process conditioned on patient-specific flow data. A KarhunenLove expansion combined with Sobol sequence sampling enables efficient exploration of the boundary condition and material parameter space for fluidsolid simulations. Gaussian process regression is then performed on a reduced-dimensional subspace obtained via principal component analysis. The resulting surrogate enables Bayesian inverse methods, including Markov chain Monte Carlo, and pose an inverse problem to reconstruct the CSF waveform at the FM. This approach demonstrates that rapid analysis of fluidsolid interactions in a realistic spinal geometry is feasible and can inform pharmacokinetics in the nervous system.

Mark Wilkinson, Aditya Shankar Paspunurwar, Hector Gomez, Ilias Bilionis
Purdue University
wilki114@purdue, apaspunu@purdue.edu, hector-

gomez@purdue.edu, ibilion@purdue.edu

MS29

Semi-Intrusive Surrogate Modelling of the Anodizing Process for Surface Morphology Control Under Uncertainty

Semi-intrusive UQ balances intrusiveness and non-intrusiveness by avoiding reformulations of governing equations while selectively replacing costly sub-models. This yields computational efficiency with no loss of model structure or fidelity, enabling applications in material science particularly corrosion models. We apply semi-intrusive UQ to study the growth of anodic aluminum oxide (AAO) films, a multistage electrochemical process progressing from barrier layer formation through pore initiation, nucleation, and steady-state self-organization. While early and final stages are relatively well characterized, transitions through pore initiation (Phase 2) and nucleation (Phase 3) remain major sources of mechanistic uncertainty. By integrating physics-based models with stochastic representations, we estimate uncertainties in nucleation kinetics and morphology under varied scenarios and spatial distributions. Semi-intrusive surrogates are built for subprocesses, enabling efficient parameter space exploration and systematic uncertainty propagation. Our results identify parameters controlling pore initiation, quantify their importance to film morphology, and provide insight into the dynamics of self-organized AAO growth. Tracking uncertainty through coupled sub-models reveals the interaction driving honeycomb structure formation. This work highlights semi-intrusive UQ ability to bridge experiments and simulation, as well as quantifying epistemic and aleatoric uncertainties.

Tamadur Albaraghtheh
Helmholtz-Zentrum Hereon
tamadur.albaraghtheh@hereon.de

Christian Feile, Mikhail Zheludkevich, Danie Höche
Helmholtz-Zentrum Hereon
Institute of Surface Science
christian.feiler@hereon.de,
mikhail.zheludkevich@hereon.de,
daniel.hoeche@hereon.de

MS29

Bayesian Optimization for Inverse Materials Design

The discovery and development of novel materials requires systematic exploration of vast design spaces. But even with powerful simulation methods like density functional theory (DFT), materials discovery remains exceptionally challenging as the vast design space makes exhaustive exploration computationally prohibitive. Thus, efficient methods for the targeted design of materials with prescribed properties, or "inverse materials design", are needed. This talk will introduce Bayesian optimization (BO) as a framework for inverse design in solid-state materials. BO is well-suited for these problems because it can identify promising candidates with relatively few expensive evaluations, quantify uncertainty to guide decision-making, and balance exploration with exploitation. I will discuss how BO can be adapted to the specific challenges of materials design, including high-dimensional search spaces, mixed variables, and competing objectives.

Nathanael Bosch
EPFL
nathanael.bosch@epfl.ch

MS29

Propagation of Quantum Scale Uncertainty Via Multivariate Conformal Prediction

In recent years, machine learned interatomic potentials (MLIPs) have become key tools for efficiently incorporating quantum scale data into a multiscale workflow. The typical representations of uncertainty of MLIPs prediction have been the posterior predictive distribution from Gaussian process regression or the standard deviation of neural network ensembles, but these indicators can exhibit unpredictable behavior on chemistry data sets. By contrast, conformal prediction provides confidence sets with coverage guarantees with minimal assumptions on the underlying data and surrogate model. Thus far, conformal prediction has been applied to quantum scale predictions component-wise to construct an predictive interval for each scalar quantity of interest (QoI). However, many downstream QoIs depend on multiple related properties, for instance both energies and forces, and on multiple molecular configurations. Gaussian process surrogates have the ability to represent these relationships and predict the covariance of vector valued QoIs for multiple configurations. This talk introduces multivariate conformal prediction for quantum chemistry and demonstrates its advantages for downstream QoI compared with component-wise conformal prediction.

Katharine Fisher
MIT
kefisher@mit.edu

MS30

State Estimation and Sensitivity Analysis of the Tropical Pacific Ocean with the Adjoint Method

Estimation and prediction of the tropical Pacific Ocean is critical to forecasting both weather and climate. This region is home to the El Nio-Southern Oscillation (ENSO), which is the dominant driver of interannual climate variability. Strong El Nios have especially severe ecosystem, human health, and economic impacts. ENSO dynamics are also notoriously challenging to model due to their asymmetric, nonlinear nature. Therefore, estimating the tropical Pacific Ocean state is a high-dimensional, nonlinear data assimilation problem, made especially challenging by limited observations. In this work we use 4D-Variational Data Assimilation to estimate the tropical Pacific Ocean state by constraining the MIT General Circulation Model (MITgcm) to observations. Using the adjoint model we pursue questions such as: 1) which observations constrain the model best, 2) what limits our ability to control the forward model run, and 3) how can we use maps of adjoint sensitivities to identify regions of high uncertainty? As an example, we will show the impact of assimilating moored velocity measurements on the representation of Equatorial phenomena like Tropical Instability Waves.

Ellen Davenport
Scripps Institution of Oceanography
edavenport@ucsd.edu

MS30

An Alternative Approach to Ensemble-based Covariance Propagation for Data Assimilation Appli-

cations

The propagation, or evolution, of the estimation error covariance is an important aspect of the statistical estimation of dynamical systems, such as data assimilation. Ensemble-based data assimilation schemes, such as the ensemble Kalman filter, approximate the evolution of the covariance through the propagation of the individual ensemble members. Thus, it is typically assumed that if the discrete state propagation and resulting mean state estimates are accurate, then the ensemble-based discrete covariance propagation will be accurate as well, apart from sampling errors due to limited ensemble size. Through a series of numerical experiments supported by analytical results, we demonstrate that this assumption is false when correlation length scales approach grid resolution. We show for states satisfying advective dynamics, that while the discrete state propagation and ensemble mean state estimates are accurate, the corresponding ensemble covariances can be remarkably inaccurate, well beyond that expected from sampling errors or typical numerical discretization errors. Motivated by these results, we propose an alternative method of discrete covariance propagation, derived from the continuum covariance dynamics, that directly addresses these errors observed during discrete covariance propagation while remaining computationally tractable in large-scale systems. Aspects of this work can then be extended to new methods of covariance modeling and covariance localization.

Shay Gilpin
University of Arizona
sgilpin@arizona.edu

MS30

Vecchia Approximation and Its Potential Applications in Geophysics

The Vecchia approximation is a computationally efficient algorithm to approximate large matrices. It doesn't even have to look at all the entries in the matrix once! Our recent work has investigated the Vecchia approximation and its mathematical optimality properties. Because the Vecchia approximation uses just a limited number of entries, it may be especially robust to the curse of dimensionality. Specifically, in geophysics it may be suitable for estimating covariance matrices where the ensemble size is many orders of magnitude smaller than the dimensionality of the state space.

Eagan Kaminetz
University of California, San Diego
ekaminetz@ucsd.edu

MS31

A Contrastive Learning Approach to Data Assimilation

Multimodal contrastive learning is a methodology for linking different data modalities, such as images and text. It is typically framed as the identification of a set of encodersone for each modalitythat align representations within a common latent space. In this presentation, we interpret contrastive learning as the optimization of encoders that define conditional probability distributions, for each modality conditioned on the other, in a way consistent with the available data. This probabilistic perspective suggests two natural generalizations of contrastive learning: (i) the introduction of novel probabilistic loss functions, and (ii) the use of alternative metrics for measuring alignment in the common latent space. We investigate these generalizations of the classical approach in the multivariate Gaussian setting by viewing latent space identification as a low-rank matrix approximation problem. The proposed framework is further studied through numerical experiments on multivariate Gaussians, the labeled MNIST dataset, and a data assimilation application in oceanography.

Ricardo Baptista
University of Toronto
r.baptista@utoronto.ca

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

Son Tran
Amazon
sontran@amazon.com

MS31

Transformers for Nonlinear Data Assimilation

Attention mechanisms, the core of transformer architectures that power modern large language models, have been widely successful at modeling nonlocal correlations in data. Recent interest in applying the transformer architecture to scientific machine learning motivates the development of rigorous mathematical analysis for this methodology. In this talk we show how the mathematical properties of the attention mechanism enable the design of transformer architectures well suited for data assimilation tasks. Namely, we demonstrate that the formulation of attention as an operator mapping between spaces of functions allows for transformer-based methodology for state estimation in nonlinear dynamical systems. On the other hand, a formulation of attention as a map between spaces of probability measures allows for methodology performing stochastic filtering. We show how this novel filtering approach presents a first step towards overcoming the Gaussian assumptions in popular model-driven data assimilation methods for nonlinear filtering.

Edoardo Calvello
California Institute of Technology
e.calvello@caltech.edu

MS31

On the Estimation Error of Optimal Transport Particle Filters

Optimal Transport Filters (OTFs) are a new family of algorithms that utilize conditional optimal transport maps for filtering of highly non-Gaussian processes. While these algorithms have demonstrated interesting practical properties their theoretical foundation is scarce. In this talk we will provide an initial attempt in this direction by outlining a series of error bounds in terms of sample complexities of a training process as well as approximation biases due to discretization. We show that a slow rate can be achieved for broad families of problems while a fast rate of convergence is contingent on theoretical constraints that are hard to satisfy and lead to interesting open questions in the theory of optimal transportation.

Bamdad Hosseini, Michele Martino
University of Washington

bamdadh@uw.edu, michem29@uw.edu

MS31

The Ensemble Score Filter for High-Dimensional Nonlinear Data Assimilation

We propose an ensemble score filter (EnSF) for solving high-dimensional nonlinear filtering problems with superior accuracy. A major drawback of existing filtering methods, e.g., particle filters or ensemble Kalman filters, is the low accuracy in handling high-dimensional and highly nonlinear problems. EnSF addresses this challenge by exploiting the score-based diffusion model, defined in a pseudo-temporal domain, to characterize the evolution of the filtering density. EnSF stores the information of the recursively updated filtering density function in the score function, instead of storing the information in a set of finite Monte Carlo samples (used in particle filters and ensemble Kalman filters). Unlike existing diffusion models that train neural networks to approximate the score function, we develop a training-free score estimation method that uses a mini-batch-based Monte Carlo estimator to directly approximate the score function at any pseudo-spatialtemporal location, which provides sufficient accuracy in solving high-dimensional nonlinear problems while also saving a tremendous amount of time spent on training neural networks. High-dimensional Lorenz-96 systems are used to demonstrate the performance of our method. EnSF provides superior performance, compared with the state-of-the-art LEnTKF, in reliably and efficiently tracking extremely high-dimensional Lorenz systems (up to 1,000,000 dimensions) with highly nonlinear observation processes.

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS32

Nested Operator Inference for Multifidelity Uncertainty Quantification in Ice Sheet Simulations

Highly accurate full-order models are often too expensive computationally to evaluate in predictive, real-time, or many-query applications. Projection-based model order reduction methods exploit the intrinsic low-dimensionality of the full-order solution manifold. These reduced-order models (ROMs) typically achieve significant computational savings while remaining physically interpretable through the governing equations. Operator Inference (OpInf) is a data-driven learning technique to construct projection-based ROMs without accessing the full-order operators. Nested OpInf exploits the inherent hierarchy within the reduced space to iteratively construct initial guesses for the OpInf learning problem that prioritize the interactions of the dominant modes. The initial guess computed for any target reduced dimension corresponds to a ROM with provably smaller or equal snapshot reconstruction error than with standard OpInf, thereby alleviating the need for careful regularization. We demonstrate the nested OpInf algorithm on a large-scale, parameterized model of the Greenland ice sheet in the context of a multifidelity uncertainty quantification study.

Nicole Aretz
The University of Texas at Austin
nicole.aretz@austin.utexas.edu

Karen E. Willcox
UT Austin
kwillcox@oden.utexas.edu

MS32

Tackling Uncertainty in Model Reduction and Scientific Machine Learning for a Catalytic CO2 Methanation Reactor

In recent years, model-based approaches have become computationally intensive, limiting their application in demanding, real-time scenarios. This highlights the need for surrogate models that are computationally efficient and capable of accurately capturing the essential dynamics of complex systems. However, for such approaches, robustness to noisy and uncertain data is required, especially in the context of processing data generated on the spot. In this study, we examine several techniques for reduced-order modeling of complex processes, including operator inference with stability guarantees and sparse identification with greedy sampling. These methods are evaluated for their ability to model the dynamic operation and control of our application, e.g., a catalytic carbon dioxide methanation reactor. We show how artificially-generated data, together with real data (measured directly from the physical entity), can be used to extract valuable information for controlling the complex processes that describe the reactor's physics under a variety of parameters and conditions. The latter category includes the excitation of the mass inlet flows using tailored inputs to mimic load variations from fluctuating hydrogen availability. By evaluating the accuracy and computational efficiency of the proposed methods, we aim to identify suitable reduced-order predictive surrogate models tailored to the challenging application under study.

Luisa Peterson
Max Planck Institute, Magdeburg, Germany
peterson@mpi-magdeburg.mpg.de

Alexander Geschke, Ion Victor Gosea, Peter Benner
Max Planck Institute Magdeburg
geschke@mpi-magdeburg.mpg.de,
gosea@mpi-magdeburg.mpg.de,
benner@mpi-magdeburg.mpg.de

Kai Sundmacher
Max Planck Institute for Dynamics of Complex Technical Sys.
sundmacher@mpi-magdeburg.mpg.de

MS32

Variationally Stable Reduced Basis Neural Operator

Minimizing PDE-based residual losses is a common way to softly enforce physical consistency in neural operators. However, the link between small residuals and small solution errors is often unclear. We show with two canonical examples, Poisson and linear elasticity, how to construct variationally correct first-order least-squares residuals that are provably equivalent to errors in PDE-compliant energy norms rather than in a generic L2 norm. This yields training objectives that directly control accuracy and naturally accommodate mixed Dirichlet/Neumann data via variational lifts. To alleviate optimization challenges and reduce training costs, we apply proper orthogonal decomposition (POD) to high-fidelity snapshots, training the neural operator to predict only the POD coefficients. This approach preserves residualerror equivalence while enabling efficient evaluation of the quadratic loss in the reduced

space. Numerical studies compare this approach with pointwise training- and evaluation-based neural operators and, importantly, show that the proposed variationally correct loss serves as a reliable a posteriori error estimator, exhibiting a tight correlation with the solution error in the appropriate PDE energy norms.

Yuan Qiu
Georgia Institute of Technology
yuan.qiu@gatech.edu

Wolfgang Dahmen
South Carolina State University
dahmen@math.sc.edu

Peng Chen
Georgia Institute of Technology
pchen402@gatech.edu

MS32

Efficient Greedy Sampling for High-dimensional Parameter Spaces

Configuration Optimization Problems (COPs), which involve minimizing a loss function over a set of discrete points $\gamma \subset P$, are common in areas like Model Order Reduction, Active Learning, and Optimal Experimental Design. While exact solutions are often infeasible, heuristic methods such as the Greedy Sampling Method (GSM) provide practical alternatives, particularly for low-dimensional cases. GSM recursively updates γ by solving a continuous optimization problem, which is typically approximated by a search over a discrete sample set $S \subset P$. However, as the dimensionality grows, the sample size suffers from the curse of dimensionality. To address this, we introduce the Polytope Division Method (PDM), a scalable greedy-type approach that adaptively partitions the parameter space and targets regions of high loss. PDM achieves linear scaling with problem dimensionality and offers an efficient solution approach for high-dimensional COPs, overcoming the limitations of traditional methods.

Karen Veroy-Grepl
TU/e Eindhoven University of Technology
k.p.veroy@tue.nl

Evie Nielen
Eindhoven University of Technology
e.t.e.nielen@tue.nl

Oliver Tse
Eindhoven University of Technology (Tu/e)
o.t.c.tse@tue.nl

MS33

Bayesian Inference of Hybrid Physics Machine Learning Models for Monitoring and Identification of Dynamical Systems

Bayesian filtering techniques are a powerful tool to monitor the behavior of dynamical systems. The Unscented Kalman Filter (UKF) is particularly well suited for damage detection, enabling both state and parameter estimation in nonlinear structural systems. However, the UKF assumes that the system dynamics are fully known, making it sensitive to model inaccuracies, unmodeled dynamics, and environmental noise. We present a hybrid framework that integrates data-driven probabilistic models (e.g., Gaussian Processes GPs or Bayesian neural networks) to model the unknown system dynamics and uncertainties without requiring explicit formulation. This hybrid framework thus aims to combine the interpretability of physics-based methods and the adaptability and robustness of data-driven approaches, while maintaining the computational efficiency of the UKF for training. We will more specifically discuss some of the challenges that arise when integrating data-driven models, which are traditionally trained off-line via optimization, within sequential learning frameworks. Also, contrary to methods that fully represent the dynamics via data-driven models without integrating known physics, the combination of both physics- and data-driven models can generate additional issues of identifiability. We will illustrate the need for robust estimation of uncertainties in this hybrid modeling context using several examples related to structural identification.

Audrey Olivier, Nihan Bilgin
University of Southern California
audreyol@usc.edu, nbilgin@usc.edu

MS33

Stochastic Sensitivity Analysis for Dynamical Systems

We derive the evolution equation for local sensitivities of statistical quantities of interest with respect to model parameters in dynamical systems with uncertain initial conditions or parameters. Our main contribution is the extension of adjoint-based a posteriori analysis for differential operators of generic dynamical systems acting on states to the Liouville operator acting on probability densities of the states. This results in theoretically rigorous estimates of sensitivity and error for a broad class of computed quantities of interest while propagating uncertainty through dynamical systems. We also derive Monte-Carlo type estimators to make these estimates computationally tractable using spatio-temporal normalizing flows and exploiting the hyperbolic nature of the Liouville equation. This work was published in the Proceeding of Machine Learning Research at Artificial Intelligence and Statistics conference in 2025.

Nishant Panda
Los Alamos National Lab
nishpan@lanl.gov

Jehanzeb H. Chaudhry
University of New Mexico
jehanzeb@unm.edu

Natalie Klein
Los Alamos National Laboratory
neklein@lanl.gov

James Carzon
Carnegie Mellon University
jcarzon@andrew.cmu.edu

Troy Butler
National Science Foundation
tdbutler@nsf.gov

MS33

Uncertainty Quantification Across the Lifecycle of

Foundation Models for Scientific Machine Learning

The recent development of data-driven foundation models for scientific machine learning (SciML) has enabled generalizable modelling of complex physical systems, but it has also amplified the challenge of understanding and managing uncertainty. Uncertainty Quantification (UQ) for foundation models must be considered at all stages of the model development lifecycle. In this work, we present a systematic perspective on UQ across four critical stages: (1) Data and Dataset Curation; (2) Tokenization and discretization; (3) Model Architecture; and (4) Training and Fine-Tuning to quantify epistemic and aleatoric uncertainties. We also perform a comparative study of SciML foundation models spanning several orders of magnitude in number of parameters (ranging from specialised architectures with a few thousand parameters to large-scale foundation models with billions of parameters). Through our study, we highlight how uncertainty evolves with model and data scale and task complexity during pretraining and finetuning. Our proposed approach will enable researchers and foundation model developers to effectively probe their development pipelines and create reliable models.

Mohd Zaki, Ponkrshnan Thiagarajan
Johns Hopkins University
mzaki4@jh.edu, ponkrshnan@jhu.edu

Souvik Chakraborty
Department of Applied Mechanics
Indian Institute of Technology Delhi
souvik@am.iitd.ac.in

Michael D. Shields
Johns Hopkins University
michael.shields@jhu.edu

MS33

Characterizing Errors in Fine-Tuned Scientific Foundation Models via Stochastic Attention

Inspired by the success of foundation models in language modeling, scientific foundation models (SciFMs) are being developed rapidly across many domains. These large-scale AI models are trained on broad scientific data and can be adapted to downstream tasks via fine-tuning. Yet for scientific deployment, we still lack practical error characterization and robust uncertainty quantification for these models. In this talk, we introduce **stochastic attention**, a mechanism that perturbs attention layers in transformer-based SciFMs to induce a controlled distribution over predictions. This provides a probabilistic reinterpretation of attention and can be viewed as a generative mechanism for modeling long-range dependencies. We use stochastic attention to quantify postfine-tuning model error with negligible runtime and storage overhead, without requiring changes to the fine-tuning procedure. We demonstrate the approach on a SciFM for weather and climate prediction, and compare it against several probabilistic fine-tuning baselines, including parametric estimations and approximate Bayesian methods.

Akash Yadav, Taiwo Adebiyi, Ruda Zhang
University of Houston
ayadav4@uh.edu, taadebiyi2@uh.edu, rudaz@uh.edu

MS34

Learning Stochastic Reduced-Order Models for Closure Terms and Model Error in High-Dimensional Plasma

High-dimensional plasma simulations, such as those performed with particle-in-cell (PIC) methods, provide critical insights into kinetic plasma behavior but are computationally prohibitive for large-scale or long-term studies. Fluid plasma models offer a more tractable alternative, yet suffer from limited physical fidelity due to inherent closure approximations and loss of key kinetic information. We presents a data-driven framework for constructing stochastic reduced-order models (ROMs) that systematically exploit the underlying low-dimensional structure of high-fidelity plasma data. By projecting rich kinetic behavior available in PIC simulations onto a lower-dimensional manifold, we hope to enable the direct learning of closure termssuch as heat fluxes, anisotropic pressure tensors, and collisional dynamicswith greater accuracy and physical consistency. We investigate both linear methods, including Dynamic Mode Decomposition (DMD), and nonlinear machine learning approaches, notably Fourier Neural Operators (FNOs), to discover accurate and interpretable representations of closure physics. Furthermore, we employ stochastic modeling strategies to rigorously quantify uncertainties and model errors introduced by dimensionality reduction and other physical assumptions. We demonstrate that by exploiting low-dimensional structure in high-fidelity PIC simulations, data-driven ROMs provide a promising path for accelerating plasma simulations and building reliable fluid closures.

Josh Eckels, Alex Gorodetsky
University of Michigan
eckelsjd@umich.edu, goroda@umich.edu

MS34

Neural Operators and Associated Uncertainty Quantification Processes

In this talk we will discuss challenges in recently developed neural operator-based models for complex structures with nonlinear dynamic behavior like offshore wind turbines. We focus on an uncertainty analysis of a virtual sensing application in which limited observations from such structures must be extrapolated to unobserved locations e.g. underwater locations.

Rad Haghi, Bipin Gaikwad, Abani Patra
Tufts University
rad.haghi@tufts.edu, bipin.gaikwad@tufts.edu, abani.patra@tufts.edu

MS34

Fokker-Planck-Based Inverse Reinforcement Learning

We introduce Fokker-Planck inverse reinforcement learning (FP-IRL), a physics-constrained IRL framework for systems governed by Fokker-Planck dynamics. Unlike standard IRL, FP-IRL infers both reward and transition functions directly from trajectory data without sampled transitions. By leveraging a conjectured equivalence between MDPs and the FP equation, we recover full MDP components via variational system identification. Experiments on synthetic and benchmark problems demonstrate accurate, interpretable, and efficient recovery of agent incentives.

Chengyang Huang
University of Michigan
chengyah@umich.edu

Siddhartha Srivastava
University of Southern California
Aerospace and Mechanical Engineering
ss79128@usc.edu

Xun Huan
University of Michigan
xhuan@umich.edu

Krishna Garikipati
USC
garikipa@usc.edu

MS35

Score-Based Generative Models for Posterior Sampling in Function Space

A central challenge in Bayesian inverse problems—particularly those governed by partial differential equations—is that the unknowns are functions living in infinite-dimensional spaces. It is therefore crucial to design inference methods that are both theoretically sound and computationally effective in arbitrarily high dimensions. Recently, score-based generative models (SGMs) have been successfully applied to Bayesian inverse problems. However, existing approaches typically assume that the posterior is supported on a finite-dimensional space, leaving the challenges of infinite dimensions to heuristics and ad hoc solutions. In this talk, we address this gap by developing a theoretical framework for applying SGMs directly to infinite-dimensional Bayesian inverse problems. We formulate SGMs at the level of function spaces and analyze their stability and convergence under refinement of the discretization. This provides a principled foundation for the use of SGMs in high- and infinite-dimensional Bayesian inference.

Lorenzo Baldassari
Universität Basel
lorenzo.baldassari@unibas.ch

Ali Siahkoohi
Georgia Institute of Technology
alisk@ucf.edu

Josselin Garnier
Ecole Polytechnique
josselin.garnier@polytechnique.edu

Knut Solna
University of California at Irvine
ksolna@math.uci.edu

Maarten V. de Hoop
Rice University
mdehoop@rice.edu

MS35

Interacting Particle Systems in Machine Learning and Inference

The convergence rate of the standard Stein Variational Gradient Descent (SVGD) is a well-known open problem. We address this gap by introducing a second-order, accelerated SVGD and proving its exponential convergence to the target distribution. Our analysis, based on a hypocoercivity framework, establishes the first general convergence rate for an SVGD-type algorithm.

Trevor Teolis
Rice University
tt111@rice.edu

MS35

An Eulerian Perspective on Straight-Line Sampling

We study dynamic measure transport for generative modeling: specifically, flows induced by stochastic processes that bridge a specified source and target distribution. The conditional expectation of the process' velocity defines an ODE whose flow map achieves the desired transport. We ask *which processes produce straight-line flows*—i.e., flows whose pointwise acceleration vanishes and thus are exactly integrable with a first-order method? We provide a concise PDE characterization of straightness as a balance between conditional acceleration and the divergence of a weighted covariance (Reynolds) tensor. Using this lens, we fully characterize affine-in-time interpolants and show that straightness occurs exactly under deterministic endpoint couplings. We then generalize this framework to stochastic processes satisfying Ito-typo SDEs, showing they obey the exact same PDE thus allowing us to study straightness for a much broader class of flows. Finally, we use our PDE framework to derive geometric conditions for straight-line inducing general processes, offering broad guidance for designing transports that are easier to integrate.

Panos Tsimpos
MIT
ptsimpos@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS35

Measure Theoretic Approaches for Uncertainty Propagation

Uncertainty is ubiquitous: both data and physical models inherently contain uncertainty. Therefore, it is crucial to identify the sources of uncertainty and control its propagation over time. In this talk, I will introduce two approaches to address this uncertainty propagation problemone for the inverse problem and one for the forward problem. The main idea is to work directly with probability measures, treating the underlying PDE as a pushforward map. In the inverse setting, we will explore various variational formulations, focusing on the characterization of minimizers and their stability. In the forward setting, we aim to propose a new approach to tackle high-dimensional uncertainties.

Li Wang
University of Minnesota
wang8818@umn.edu

MS36

Operator Learning with Gaussian Processes: From Handling Point Clouds to Addressing Scalability and UQ

Gaussian processes (GPs) are powerful machine learning models that have been successfully used in many applications that involve clustering, density estimation, system identification, regression, and classification. The major-

ity of recent works on GPs are aimed at increasing their scalability to high dimensions and/or large datasets but GPs are getting increasing attention in operator learning and solving forward/inverse problems involving PDEs. In this talk, we will focus on some of the fundamental limitations of kernels in the context of scientific machine learning. By addressing these limitations, we demonstrate that GPs can provide computationally efficient and competitive performance in critical applications such as topology optimization and operator learning.

Ramin Bostanabad
University of California, Irvine
raminb@uci.edu

MS36

Sparse Gaussian Process Operator for Uncertainty Quantification

Operator learning offers a powerful paradigm for solving parametric partial differential equations (PDEs), but scaling probabilistic neural operators such as the recently proposed Gaussian Processes Operators (GPOs) to high-dimensional, data-intensive regimes remains a significant challenge. In this work, we introduce a novel, scalable GPO, which capitalizes on sparsity, locality, and structural information through judicious kernel design. Addressing the fundamental limitation of cubic computational complexity, our method leverages nearest-neighbor-based local kernel approximations in the spatial domain, sparse kernel approximation in the parameter space, and structured Kronecker factorizations to enable tractable inference on large-scale datasets and high-dimensional input. While local approximations often introduce accuracy trade-offs due to limited kernel interactions, we overcome this by embedding operator-aware kernel structures and employing expressive, task-informed mean functions derived from neural operator architectures. Through extensive evaluations on a broad class of nonlinear PDEs, we demonstrate that our framework consistently achieves high accuracy across varying discretization scales.

Souvik Chakraborty
Department of Applied Mechanics
Indian Institute of Technology Delhi
souvik@am.iitd.ac.in

MS36

Deep Operator Gaussian Process

We present a two-stage framework, Deep Operator Gaussian Processes (DeepOp-GP), for uncertainty quantification in operator learning. A high-dimensional data GP first models inputoutput mappings of operators, providing predictive means and covariances. A function GP with deep kernel learning then conditions on these uncertain outputs, yielding closed-form predictive variance that combines both sources of uncertainty. We demonstrate the method on toy operators, nonlinear PDEs (Burgers, reactiondiffusion), and climate modeling data, showing improved accuracy and calibrated uncertainty estimates for high-dimensional scientific problems.

Jaeyong Lee
Chung-Ang University, Korea
jaeyong@cau.ac.kr

MS36

Epistemic Physics-Informed Neural Networks

Physics-informed neural networks (PINNs) have demonstrated promise as a framework for solving forward and inverse problems involving partial differential equations. Despite recent progress in the field, it remains challenging to quantify uncertainty in these networks. While approaches such as Bayesian PINNs (B-PINNs) provide a principled approach to capturing uncertainty through Bayesian inference, they can be computationally expensive for large-scale applications. We propose Epistemic Physics-Informed Neural Networks (E-PINNs), a framework that leverages a small network, the *epinet*, to efficiently quantify uncertainty in PINNs. The proposed approach works as an add-on to existing, pre-trained PINNs with a small computational overhead. We demonstrate the applicability of the proposed framework in various test cases and compare the results with B-PINNs using Hamiltonian Monte Carlo (HMC) posterior estimation and dropout-equipped PINNs (Dropout-PINNs). Our experiments show that E-PINNs provide similar coverage to B-PINNs, with often comparable sharpness, while being computationally more efficient. This observation, combined with E-PINNs' more consistent uncertainty estimates and better calibration compared to Dropout-PINNs for the examples presented, indicates that E-PINNs offer a promising approach in terms of accuracy-efficiency trade-off.

Panos Stinis
Pacific Northwest National Laboratory
panagiotis.stinis@pnnl.gov

MS37

Metropolis-Adjusted Subdifferential Langevin Algorithm

The Metropolis-Adjusted Langevin Algorithm (MALA) is a widely used Markov Chain Monte Carlo (MCMC) method for sampling from high-dimensional distributions. However, MALA relies on differentiability assumptions that restrict its applicability. In this paper, we introduce the Metropolis-Adjusted Subdifferential Langevin Algorithm (MASLA), a generalization of MALA that extends its applicability to distributions whose log-densities are locally Lipschitz, generally non-differentiable, and nonconvex. We evaluate the performance of MASLA by comparing it with other sampling algorithms in settings where they are applicable. Our results demonstrate the effectiveness of MASLA in handling a broader class of distributions while maintaining computational efficiency.

Ning Ning
Texas A&M University
patning@tamu.edu

MS37

Optimal Transport-Based Generative Models for Bayesian Posterior Sampling

We investigate the problem of sampling from posterior distributions with intractable normalizing constants in Bayesian inference. Our solution is a new generative modeling approach based on optimal transport (OT) that learns a deterministic map from a reference distribution to the target posterior through constrained optimization.

The method uses structural constraints from OT theory to ensure uniqueness of the solution and allows efficient generation of many independent, high-quality posterior samples. The framework supports both continuous and mixed discrete-continuous parameter spaces, with specific adaptations for latent variable models and near-Gaussian posteriors. Beyond computational benefits, it also enables new inferential tools based on OT-derived multivariate ranks and quantiles for Bayesian exploratory analysis and visualization. We demonstrate the effectiveness of our approach through multiple simulation studies and a real-world data analysis.

Yuexi Wang
University of Illinois Urbana-Champaign
yxwang99@illinois.edu

MS38

Continuous Adaptive Nonlinear Optimal Design

Identifying informative measurements is a central challenge in computational inverse problems. This task is often framed as an optimal experimental design (OED) problem. In this talk, we explore OED in the context of a continuously-indexed design space, as opposed to traditional approaches that choose measurements from a finite set. This continuous design formulation better reflects practical scenarios where measurements are seamlessly distributed over spatial or temporal domains. To address the associated computational and optimization challenges, we leverage gradient flow methods, incorporated by adaptive algorithms for models with nonlinear parameter-to-observation (PtO) maps. The proposed algorithm not only captures the most informative measurements but also adaptively leads to qualified reconstruction of inverse problems. This novel OED framework has shown effectiveness across a range of physical models, from quantum systems to dynamical models.

Ruhui Jin
University of Wisconsin
rjin18@jh.edu

MS38

Efficient Optimization of Expensive Black-Box Simulators Via Marginal Means

With advances in scientific computing, computer experiments are increasingly used for optimizing complex systems. However, for modern applications, e.g., the optimization of nuclear physics detectors, each experiment run can require hundreds of CPU hours, making the optimization of its black-box simulator f over a high-dimensional space a challenging task. Given limited runs at inputs $\mathbf{x}_1, \cdots, \mathbf{x}_n$, the best solution from these evaluated inputs can be far from optimal, particularly as dimensionality increases. Existing black-box methods, however, largely employ this "pick-the-winner" (PW) solution, which leads to mediocre optimization performance. To address this, we propose a new Black-box Optimization via Marginal Means (BOMM) approach. The key idea is a new estimator of a global optimizer $\mathbf{x}^*$ that leverages the so-called marginal mean functions, which can be efficiently inferred with limited runs in high dimensions. Under mild conditions, we prove that the BOMM estimator not only is consistent for optimization, but also has an optimization rate that tempers the "curse-of-dimensionality" faced by existing methods, thus enabling better performance as dimensionality increases. We present a practical framework for implementing BOMM using the transformed Gaussian process surrogate model.

Hwanwoo Kim
Duke University
ghksdn1227@gmail.com

MS38

Reducing the Dimension of Single-Cell Data: Multi-Group Structures and Periodicity

Insingle-celldataanalysis, recent progress on dimension reduction has focused on leveraging group structures and periodicity. Advanced statistical methods are utilized to identify unique structures in distinct groups, aiding in distinguishing between healthy and diseased cells. This is key for uncovering complex cellular behaviors and understanding disease mechanisms. The unique feature of periodicity in the cell data enhances the interpretability of visualizations in complex biologicaldata. These advancements markstrides in statistical methodologies forsingle-celldataanalysis, emphasizing the need for precision and sophisticated algorithms, ultimately leading to our new methodologies.

Hengrui Luo
Lawrence Berkeley National Laboratory
hengruiluo@gmail.com

MS38

Sanvi: A Fast Spectral-Assisted Network Variational Inference Method with An Extended Surrogate Likelihood Function

Bayesian inference has been broadly applied to statistical network analysis, but suffers from the expensive computational costs due to the nature of Markov chain Monte Carlo sampling algorithms. This paper proposes a novel and computationally efficient Spectral-Assisted Network Variational Inference (SANVI) method within the framework of the generalized random dot product graph. The key idea is a cleverly designed extended surrogate likelihood function that enjoys two convenient features. Firstly, it decouples the generalized inner product of latent positions in the random graph model. Secondly, it relaxes the complicated domain of the original likelihood function to the entire Euclidean space. Leveraging these features, we design a computationally efficient Gaussian variational inference algorithm via stochastic gradient descent. Furthermore, we show the asymptotic efficiency of the maximum extended surrogate likelihood estimator and the Bernstein-von Mises limit of the variational posterior distribution. Through extensive numerical studies, we demonstrate the usefulness of the proposed SANVI algorithm compared to the classical Markov chain Monte Carlo algorithm, including comparable estimation accuracy for the latent positions and less computational costs.

Fangzheng Xie
Indianna University
Department of Statistics
fxie@iu.edu

MS39

Learning Enhanced Ensemble Filters

The filtering distribution in hidden Markov models evolves according to the law of a mean-field model in state-observation space. The ensemble Kalman filter (EnKF)

approximates this mean-field model with an ensemble of interacting particles, employing a Gaussian ansatz for the joint distribution of the state and observation at each observation time. These methods are robust, but the Gaussian ansatz limits accuracy. This shortcoming is addressed by approximating the mean-field evolution using a novel form of neural operator taking probability distributions as input: a measure neural mapping (MNM). A MNM is used to design a novel approach to filtering, the MNM-enhanced ensemble filter (MNMEF), which is defined in both the mean-field limit and for interacting ensemble particle approximations. The ensemble approach uses empirical measures as input to the MNM and is implemented using the set transformer, which is invariant to ensemble permutation and allows for different ensemble sizes. The derivation of methods from a mean-field formulation allows a single parameterization of the algorithm to be deployed at different ensemble sizes. In practice fine-tuning of a small number of parameters, for specific ensemble sizes, further enhances the accuracy of the scheme. The promise of the approach is demonstrated by its superior root mean-square-error performance relative to leading methods in filtering the Lorenz 96 and Kuramoto-Sivashinsky models.

Bohan Chen
California Institute of Technology
bhchen@caltech.edu

MS39

Response Theory Via Score Modeling

Predicting how small forcings reshape the full probability distribution of complex systems is essential for design, risk, and science. I will present a data-driven framework that combines the Generalized FluctuationDissipation Theorem (GFDT) with score-based generative modeling to forecast distributional responsesmean, variance, skewness, and kurtosisin nonlinear, non-Gaussian dynamics. The key is to estimate the score (the gradient of the log steady-state density) directly from data: for low-dimensional effective dynamics we use a clustering-based estimator (KGMM), while for high-dimensional fields we train a denoising U-Net. Across a hierarchy of testbedsscalar and multiscale reduced-order models, a barotropic atmospheric regime model, and a 2-D NavierStokes turbulence example with localized forcingthe method reproduces ground-truth responses and markedly outperforms Gaussian approximations, including in the tails relevant for extremes.

Ludovico T. Giorgini
Nordita
Stockholm University
giorgini.ludovico@gmail.com

MS39

Data Assimilation Model for Multiscale Probability Distributions of Turbulent Systems

We present a new strategy for the statistical forecasts and uncertainty quantification of multiscale nonlinear turbulent systems involving non-Gaussian probability distributions. The capability of using reduced-order models to capture key statistical features is investigated. A closed stochastic-statistical modeling framework is proposed using a high-order statistical closure enabling accurate prediction of leading-order statistical moments and probability density functions in multiscale complex turbulent systems. A new efficient ensemble forecast algorithm is developed dealing with the nonlinear multiscale coupling mechanism as a characteristic feature in high-dimensional turbulent systems. Effective nonlinear ensemble filters are developed based on the particular nonlinear coupling structures of the explicit stochastic and statistical equations, which satisfy an infinite-dimensional Kalman-Bucy filter with conditional Gaussian dynamics. It is demonstrated that crucial principal statistical quantities in the most important large scales can be captured efficiently with accuracy using the new reduced-order model in various dynamical regimes of the flow field with distinct statistical structures.

Di Qi
Purdue University
qidi@purdue.edu

MS39

Observation Network Impacts of the Ensemble Score Filter in the Surface Quasi-geostrophic System

Recent advances in data assimilation (DA) have focused on devel- oping more flexible approaches that can better accommodate nonlinearities in models and observations. However, it remains unclear how the performance of these advanced methods depends on the observation network characteristics. In this study, we present initial experiments with the surface quasi-geostrophic model, in which we compare a recently developed AI-based ensemble filter with the standard Local Ensemble Transform Kalman Filter (LETKF). Our results show that the analysis solutions respond differently to the number, spatial distribution, and nonlinear fraction of assimilated observations. We also find notable changes in the multiscale characteristics of the analysis errors. Given that standard DA techniques will be eventually replaced by more advanced methods, we hope this study sets the ground for future efforts to reassess the value of Earth observation systems in the context of newly emerging algorithms.

Zixiang Xiong
Department of Mathematics
Florida State University
zx18@fsu.edu

MS40

Large-Dimensional Shapley Effects' Estimation for Reliability-oriented Sensitivity Analysis

In the study of physical systems, reliability analysis aims at estimating failure probabilities and identifying the input variables most responsible for failures. Reliability-oriented sensitivity analysis complements probability estimation by quantifying the contribution of each input to the occurrence of rare events. When dealing with models involving many correlated inputs, classical Sobol indices are not appropriate, while Shapley effects provide a consistent framework. However, existing estimation procedures have mainly been restricted to low-dimensional models or require prohibitive computational budgets, especially when failure probabilities are very small. We propose a novel methodology for estimating reliability-oriented Shapley effects in higher-dimensional models by leveraging Normalizing Flows. Their ability to approximate high-dimensional conditional densities, combined with an efficient aggregation strategy, enables accurate estimation of Shapley effects for correlated inputs without additional simulations beyond those required for failure probability estimation. The proposed approach is validated on numerical case stud-

ies with more than 20 input variables, illustrating its effectiveness and scalability in high-dimensional and correlated settings.

Lucas Monteiro, Jérôme Morio
ONERA
lucas.monteiro@onera.fr, jerome.morio@onera.fr

Julien Demange-Chryst
ONERA
Institut Mathématiques de Toulouse
julien.demange-chryst@onera.fr

François Bachoc
Laboratoire Paul Painlevé, Université de Lille
francois.bachoc@univ-lille.fr

MS40

Conditional Feature Importance Revisited: Double Robustness, Efficiency and Inference

In this presentation, we begin by introducing a Generalized Total Sobol Index for a studied feature j. This index comes from Sensitivity Analysis and can be adapted to Feature Importance. It provides an interpretable measure that helps us carry out Conditional Independence Testing with statistical control. We then discuss three model-agnostic estimation procedures, meaning they can be applied to any black-box model: **1.LeaveOneCovariateOut(LOCO)** : compares the performance of the original model with that of a refitted version excluding feature j. **2.ConditionalFeatureImportance(CFI)** : compares the performance of the same black-box model when the j-th coordinate is conditionally perturbed against the original input. We highlight its performance in variable selection through a *double robustness* property: identifying null covariates requires either a well-specified conditional sampler or a well-specified model. **3.Sobol − CFI** : a generalization of CFI based on marginalization, i.e., the performance decay of the original model compared with its conditionally averaged version. For this procedure, we correct the bias and present a nonparametric efficiency result. Finally, we discuss type-I error control for the presented procedures and provide numerical experiments illustrating the estimation findings. The results show that Sobol-CFI retains the double robustness property in practice.

Angel D. Reyero Lobo
INRIA Paris Saclay
angel.reyero-lobo@inria.fr

Bertrand Thirion
Inria Paris Saclay
bertrand.thirion@inria.fr

Pierre Neuvial
Institut de Mathématiques de Toulouse
pierre.neuvial@math.univ-toulouse.fr

MS41

Rockafellian Relaxation for PDE Constrained Optimization with Distributional Ambiguity.

Stochastic optimization problems are generally known to be ill-conditioned to the form of the underlying uncertainty. A framework is introduced for optimal control problems with partial differential equations as constraints that is robust to inaccuracies in the precise form of the problem uncertainty. The framework is based on problem relaxation and involves optimizing a bi-variate, "Rockafellian" objective functional that features both a standard control variable and an additional perturbation variable that handles the distributional ambiguity. In the presence of distributional corruption, the Rockafellian objective functionals are shown in the appropriate settings to Γ-converge to uncorrupted objective functionals in the limit of vanishing corruption. Numerical examples illustrate the frameworks utility for outlier detection and removal and for variance reduction.

Harbir Antil
George Mason University
hantil@gmu.edu

MS41

On Invariant-Domain Preserving Uncertainty Quantification for Conservation Laws

Conservation or balance laws, partial differential equations for which the evolved state quantities are conserved, model physical phenomena spanning orders of magnitude in spatial and temporal scales. For these problems, numerical and physical complexity mandate intricate schemes to obtain satisfactory capability. Though challenging even in a deterministic setting, as a result of their practical significance, uncertainties permeate these studies: initial conditions, boundary conditions, material characteristics, and equations of state, among other properties and conditions, all contribute sources of uncertainty that preclude mere point estimates in robust, predictive computations. In this talk, we consider invariant-domain preserving uncertainty quantification based on a performant extension of recent advancements in stochastic active discretizations and stochastic finite volume methods for conservation laws. To demonstrate how these advancements enable enhanced propagation of uncertainty, we consider a series of benchmark examples based on Euler's equations.

Jake J. Harmon
Los Alamos National Laboratory
harmon@lanl.gov

Jerett Cherry
Colorado State University
Los Alamos National Laboratory
cherry@lanl.gov

Eric Tovar
Los Alamos National Laboratory
T-5 Applied Mathematics and Plasma Physics
ericjtovar@lanl.gov

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

MS41

Novel Method for Uncertainty Quantification in Hyperbolic Problems Using Linear Programming

We present an alternative approach to treat random hyperbolic equations numerically. This approach does not suffer from the loss of hyperbolicity and it allows to incorporate existing methods, like stochastic collocation methods. The approach relies on the usage of Young measures that in-

troduce an inherent linear structure on the problem. The latter is then discretized and it is included in a classical high-order central Upwind scheme. Numerical results are presented.

Michael Herty
RWTH Aachen University
herty@igpm.rwth-aachen.de

MS41

Uncertainty Reduction in Digital Twins Simulations

Creating reliable digital twins is often challenged by significant model uncertainty, including unknown governing equations and internal parameters, unobserved states, and even unspecified dimensionality. This lack of knowledge presents a fundamental barrier to traditional modeling and parameter inference techniques, making robust prediction a major challenge. We introduce a Flow Map Learning (FML) framework that bypasses these challenges by constructing a predictive model directly from observational data. Our approach uses a deep neural network to learn the evolution operator, i.e., flow map, which maps a finite history of observed states to the system's future state. The core insight is that the observational history itself implicitly encodes the effects of the unknown parameters and unobserved dynamics. By conditioning on this history, the model can accurately determine the system's trajectory, thereby resolving the core challenge of forecasting under uncertainty. The trained FML model provides a compact and efficient surrogate representation of the unknown system, enabling accurate long-horizon predictions of quantities of interest (QoI) while naturally accommodating various uncertainty sources. We demonstrate the effectiveness of our approach through several numerical examples, showcasing its capability to produce reliable predictions and reduce uncertainty in digital twin applications.

Dongbin Xiu, Xinyu Liu
Ohio State University
xiu.16@osu.edu, liu.12165@osu.edu

Zhongshu Xu
The Ohio State University
xu.4202@osu.edu

Xiaoting Han
Ohio State University
han.2051@osu.edu

MS43

Subspace Accelerated Measure Transport for Fast and Scalable Sequential Experimental Design

We propose a novel approach for sequential optimal experimental design (sOED) for Bayesian inverse problems involving expensive models with large-dimensional unknown parameters. The focus of this work is on designs that maximize the expected information gain (EIG) from prior to posterior, which is a computationally challenging task in the non-Gaussian setting. This challenge is amplified in sOED, as the incremental expected information gain (iEIG) must be approximated multiple times in distinct stages, with both prior and posterior distributions often being intractable. To address this, we derive a derivative-based upper bound for the iEIG, which not only guides design placement but also enables the construction of projectors onto likelihood-informed subspaces, facilitating parameter dimension reduction. By combining this approach with conditional measure transport maps for the sequence of posteriors, we develop a unified framework for sOED, together with amortized inference, scalable to high- and infinite-dimensional problems. Numerical experiments for two inverse problems governed by partial differential equations (PDEs) demonstrate the effectiveness of designs that maximize our proposed upper bound.

Robert Scheichl
Ruprecht-Karls University Heidelberg
r.scheichl@uni-heidelberg.de

Karina Koval
Heidelberg University
karina.koval@iwr.uni-heidelberg.de

Tiangang Cui
Monash University
tiangang.cui@sydney.edu.au

Roland Herzog
Heidelberg University
roland.herzog@iwr.uni-heidelberg.de

MS44

Distributional Reinforcement Learning for Sequential Experimental Design under Uncertainty

Distributional reinforcement learning has shown significant promise for robust decision-making under uncertainty. We introduce a distributional reinforcement learning framework for sequential experimental design, formulated as a finite-horizon partially observable Markov decision process. By learning parameterized return distributions, our method explicitly captures uncertainty in experimental outcomes and supports adaptive, risk-aware policies that move beyond greedy or batch strategies. We demonstrate the approach by using a time-dependent linear model and an advectiondiffusion system, highlighting its effectiveness for guiding experiment selection in complex, uncertain environments.

Sahil Bhola
University of Michigan, Ann Arbor
sbhola@umich.edu

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

MS44

Wildfire Perimeter Estimation from Emergency Call Data

Wildfires are increasingly reaching populated areas, heightening the need for accurate forecasts. Because real-time observations are often limited, predictions first require estimating the fire perimeter from sparse data. Among existing state reconstruction methods, Discrete Empirical Interpolation Method (DEIM) is notable for its computational efficiency and interpretability. However, DEIM may return unphysical values if the true state variable possesses upper or lower bound constraints. More specifically, DEIM estimates may contain values outside the physical bounds, such as negative mass density. In this talk, we present Constrained DEIM (C-DEIM), a method designed to ac-

curately reconstruct states with known upper and lower bounds. C-DEIM employs soft constraints to enforce these bounds, generating realistic reconstructions suitable for downstream tasks. We demonstrate the effectiveness of C-DEIM by reconstructing and forecasting the 2023 Maui wildfire perimeter from emergency call data.

Louisa Ebby, Mohammad Farazmand
North Carolina State University
lbebby@ncsu.edu, farazmand@ncsu.edu

MS44

Step-Dad: Semi-Amortized Policy-Based Bayesian Experimental Design

We develop a semi-amortized, policy-based, approach to Bayesian experimental design (BED) called Stepwise Deep Adaptive Design (Step-DAD). Like existing, fully amortized, policy-based BED approaches, Step-DAD trains a design policy upfront before the experiment. However, rather than keeping this policy fixed, Step-DAD periodically updates it as data is gathered, refining it to the particular experimental instance. This test-time adaptation improves both the flexibility and the robustness of the design strategy compared with existing approaches. Empirically, Step-DAD consistently demonstrates superior decision-making and robustness compared with current state-of-the-art BED methods.

Marcel Hedman
Jesus College
University of Oxford
marcel.hedman@jesus.ox.ac.uk

Desi R. Ivanova
Oxford
desi.ivanova@stats.ox.ac.uk

Tom Rainforth
University of Oxford
rainforth@stats.ox.ac.uk

MS45

Control of Magnetized Plasmas by External Electric Fields

The stability of magnetized plasmas is highly sensitive to model assumptions and perturbations, making quantitative analysis and control a challenging task. We investigate a control problem for uniformly magnetized plasmas with external electric fields. Linear analysis around general equilibria is carried out using a LaplaceFourier transform framework, and the associated Penrose-type stability condition for Bernstein modes is derived and numerically verified. Building on this characterization, we introduce a general control strategy that recovers the free-streaming solution as a benchmark case. Several case studies, including Gaussian equilibria and the DoryGuestHarris instability, illustrate how the proposed control strategy can robustly suppress instabilities. These results highlight a pathway for including the discussion of UQ in inverse and control problems for kinetic equations. There are interesting directions to extend this problem by bringing data-driven approaches and machine learning into the picture.

Peiyi Chen, Rogerio Jorge, Qin Li, Yukun Yue
University of Wisconsin-Madison
pchen345@wisc.edu, rogerio.jorge@wisc.edu, qinli@math.wisc.edu, yyue24@wisc.edu

MS45

Particle Methods for Plasma Equations with Uncertainties

Plasma models are studied at the kinetic level by the Landau equation and have gained a lot of interest due to the important applications related to fusion reactors and ongoing projects. Moreover, the importance of considering uncertainties into the PDEs is growing, partly due to the increased presence of real-world data. The construction of numerical methods for such equation with random inputs is a challenging problem, thanks to the high dimensional structure of the equation, involving both phase space variables and stochastic parameters, and the formation of multiscale structures that must be captured by the numerical schemes. Furthermore, a numerical solver must be able to preserve the structural physical properties such as the non-negativity of the distribution function, the main conservations of invariant quantities, the entropy dissipation, and the equilibrium states. Besides, the regularity of the solution with respect to the random parameter is crucial to derive efficient numerical methods, and must be investigated. Recently, a new class of numerical methods that combine a particle-based approximation of the distribution function in the phase space together with a stochastic Galerkin expansion of the particles in the random space has been proposed. These methods are spectrally accurate in the space of the random parameters, thanks to the sG formulation, and conserve the structural properties of the equation, because of the particle nature.

Andrea Medaglia
University of Oxford
andrea.medaglia@unife.it

MS45

Convergence of Fourier-Galerkin Spectral Method for the Boltzmann Equation with Uncertainties

In this talk, we will present the results about the Boltzmann equation with uncertainties and prove that the spectral convergence of the semi-discretized numerical system holds in a combined velocity and random space, where the Fourier spectral method is applied for approximation in the velocity space, whereas the generalized polynomial chaos (gPC)-based stochastic Galerkin (SG) method is employed to discretize the random variable. Our proof is based on a delicate energy estimate for showing the well-posedness of the numerical solution as well as a rigorous control of its negative part in our well-designed functional space that involves high-order derivatives of both the velocity and random variables.

Kunlun Qi
Simons Laufer Mathematical Sciences Institute (SLMath)
Michigan State University
kunlunqi.math@gmail.com

Liu Liu
Chinese University of Hong Kong
liuliu@cuhk.edu.hk

MS46

Interatomic Potential Uncertainty Estimation and

Propagation Via Embedded Model Error

Novel materials capable of withstanding the extreme environments in fusion energy devices are often explored via molecular dynamics (MD) simulations, enabling researchers to probe material properties and behavior under representative loading conditions. In the interest of computational efficiency, such MD simulations often involve a choice of interatomic potential (IAP) tuned to match data from more expensive, higher-fidelity approaches (i.e. density functional theory), though this choice inevitably introduces uncertainty. The present work aims to estimate the uncertainty associated with two IAPs: the Spectral Neighbor Analysis Potential (SNAP) and Atomic Cluster Expansion (ACE). Uncertainty associated with SNAP is estimated via an embedded model error approach, in which uncertainty is embedded into the model form by prescribing a distribution to the model coefficients as well as the moments that parameterize that distribution. The maximum a posteriori (MAP) estimate of the model covariance is obtained from the resulting Bayesian inverse problem, accounting for model form in uncertainty estimates in derived energies, forces, and stresses across multiple materials. The embedded model error approach is compared to traditional Bayesian linear regression. Further efforts explore the construction of a polynomial chaos surrogate for more efficient forward propagation of uncertainties in elastic properties, binding energies, and defect properties obtained with an ACE potential.

Cristian Lacey, Mary Alice Cusentino
Sandia National Laboratories
celacey@sandia.gov, mcusent@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

MS46

Dimension Reduction on Hilbert Spaces of Velocity Fields of SDEs

In a controlled stochastic differential equation (SDE), the control is a velocity field which serves to guide the solution of the SDE toward a chosen terminal time-marginal probability measure, or alternatively, to represent uncertainty in the drift coefficient of the SDE. Many problems involve optimization or sampling over a function space of controls, where either task proves challenging when the control is defined in infinite dimensions. To address this challenge, we take a dimension reduction approach of solving for an optimal rank-r subspace of a Hilbert space of permissible controls and use it to construct efficient ridge approximations of target functionals of the control. We articulate sufficient conditions on the Hilbert space of controls such that the corresponding controlled SDEs have weak solutions and that one can specify a valid Cameron-Martin space embedded in the Hilbert space. Finally, we evaluate the advantages of our dimension reduction approach for two applications, 1) minimization of the objective functional of a stochastic optimal control problem, and 2) forward uncertainty quantification of functionals which represent reaction rates in molecular dynamical systems.

Joanna Zou
Massachusetts Institute of Technology
jjzou@mit.edu

Han Cheng Lie
University of Potsdam
Institute of Mathematics
hanlie@uni-potsdam.de

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS47

Simulation-based Uncertainty Quantification for Remote Sensing Inverse Problems

Remote sensing data provide a vast trove of information about Earth systems. These data are inferred from electromagnetic spectra observed by Earth-orbiting instruments using computational algorithms that approximately solve the following inverse problem. Given a spectrum, find the best estimate of the underlying true geophysical state that gave rise to it. The algorithms typically start with an initial guess for the true state, and iteratively try to minimize cost functions based on differences between predicted spectra from a deterministic forward model, and spectra actually observed. This set-up incorporates many important aspects related to uncertainty quantification for complex physical-statistical systems including intractable likelihoods, unknown physics, parameter calibration, computational demands, etc. The remote sensing problem is also inherently spatial, which is both a curse, in that it usually involves massive amounts of data, and a blessing in that it makes dimension reduction possible if we can learn and quantify spatial structure. In this talk I will discuss recent work that demonstrates how we leverage simulation-based inference, spatial statistical modeling with Gaussian processes, and conditional spatial simulation to obtain robust uncertainty estimates for the output of a remote sensing data processing pipeline. If time permits, I will also show results of some preliminary work to quantify and incorporate model discrepancy.

Amy Braverman
Jet Propulsion Laboratory
California Institute of Technology
Amy.Braverman@jpl.nasa.gov

MS47

Accelerations to Climate Model Development Using An Ensemble Kalman Inversion Toolkit

Ensemble Kalman Processes (EKP) can be applied to calibrate model parameters of expensive computer codes from indirect statistical data. Ensembles of simulations at different model parameter configurations are iteratively updated to fit a fixed indirect statistical observation using Kalman-type updates, as seen with the ensemble Kalman filter in data assimilation (DA) literature. Viewing EKP as applications of DA methodology to inverse problems has lead to enhancements such as localization (of covariances/observations/domain) to increase scalability, and inflation to help with model error. Meanwhile, viewing EKP as an iterative black-box optimizer yields enhancements such as adaptive learning rates and momentum acceleration to prevent stagnation. We showcase how this workhorse can provide a variety of improvements to Climate model development for the CliMA (Climate Modeling Alliance) model. We showcase its use to adapt train-

ing point selection for accelerated sampling algorithms for parametric uncertainty (Calibrate Emulate Sample). We also showcase how it enables online learning for model-error diagnosis and improvement of closures for parameterizations (Calibration-assisted Development).

Oliver Dunbar, Tapio Schneider
California Institute of Technology
odunbar@caltech.edu, tapio@caltech.edu

Eviatar Bach
California Institute of Technology, Pasadena, California
eviatarbach@protonmail.com

Costa Christopolous
California Institute of Technology
cchristo@caltech.edu

MS47

Improving Uncertainty Quantification in Atmospheric Data Assimilation with Machine Learning

Atmospheric data assimilation (DA) integrates observations with model forecasts to produce optimized atmospheric states, which is critical for stable weather forecasting and reliable climate research. Modern DA is posed in a Bayesian framework that combines a model-based prior (background) with an observation likelihood. However, the high dimensionality of the atmospheric state and evolving cross-variable dependencies make estimating uncertainty in the background-error prior particularly challenging. This study leverages machine learning to address these challenges.First, we propose Latent Data Assimilation (LDA), which performs variational DA in a latent space learned from multivariate global atmospheric data via an autoencoder. The autoencoder decorrelates atmospheric variables and captures nonlinear physical relationships, enabling LDA to produce balanced analyses even with a diagonal error covariance matrix. Second, we introduce Diffusion-DA, a score-based framework that learns a generative, non-Gaussian background prior and incorporates observations to guide denoising, yielding approximate posterior samples. Beyond this, we employ preference-based reinforcement learning to calibrate the background priors uncertainty, promoting analyses that align with DA objectives such as physical consistency and improved forecast performance.

Hang Fan
Columbia University
hf2526@columbia.edu

MS47

Quantifying and Propagating Uncertainty in Cloud and Precipitation Processes in Weather and Climate Simulations Via Deterministic and Stochastic Methods

Cloud processes are a leading source of uncertainty in climate projections owing to the difficulty in representing the microscale processes that produce clouds, rain, and snow with the limited degrees of freedom available in Earth System Models (ESMs). I will present recent work addressing parametric uncertainties, quantifying them in both high resolution and global models using ML-assisted Bayesian inference. This ability is largely novel, and allows for uncertainty in climate projections to be provided, facilitating risk assessment, as well as allowing for quantification of observational information content. However, there are are also structural sources of uncertainty in representation of clouds in ESMs – these require a different approach. Machine learning has provides some benefits, with some difficulties regarding implementation. However, more fundamentally there is strong reason to believe that many uncertainties, in particular structural uncertainties, should be represented *stochastically, with spatiotemporal covariability that matches natural variations. I will present preliminary work in this direction, and discuss future steps.*

Marcus Van Lier-Walqui
Center for Climate Systems Research
Columbia University/NASA GISS
mv2525@columbia.edu

Kaitlyn Loftus
Columbia University
kaitlyn.loftus@columbia.edu

Hugh Morrison
National Center for Atmospheric Research
morrison@ucar.edu

Adele Igel
University of California - Davis
aigel@ucdavis.edu

Mikael Witte
Naval Postgraduate School
mikael.witte@nps.edu

Arthur Hu
Columbia University
zh2611@columbia.edu

MS48

Conditional Flow-Matching for Data Assimilation

Traditional data assimilation techniques face significant challenges in high-dimensional problems where process dynamics and observation operators are nonlinear, and where both process and observation noise are non-Gaussian and non-additive. In this talk, we will explore the use of flow-matching algorithms to address these challenges. Conditional flow-matching algorithms enable the transport of samples from the prior distribution to the posterior distribution using knowledge of samples from the joint distribution. We will demonstrate how such algorithms can be applied to approximate the update step in the Bayes filter for data assimilation. Furthermore, we will show how they facilitate the use of black-box forward dynamics and observation models, thereby accommodating a broad class of stochastic models for both the forward process and the observation operator. Finally, we will discuss the computational costs associated with training and applying these models, as well as their application in solving challenging data assimilation problems.

Brianna Binder
University of Southern Carolina
bjbinder@usc.edu

Agnimitra Dasgupta
Department of Aerospace and Mechanical Engineering
University of Southern California
adasgupt@usc.edu

Assad A. Oberai
University of Southern California

aoberai@usc.edu

MS48

Improving the Fidelity of PDE Data with Diffusion Models

Solving partial differential equations (PDEs) on fine spatio-temporal scales for high-fidelity solutions is critical for scientific discoveries, but can be prohibitively expensive. As an alternative, super-resolution techniques enable the enhancement of low-fidelity, low-resolution simulations. In this talk, I will present a family of diffusion-based super-resolution methods that leverage score-based generative modeling to infer sub-grid information while respecting the governing physics. Through extensive numerical experiments, we demonstrate that these approaches retain the accuracy of full-resolution solvers while significantly reducing computational expense.

Yulong Lu
University of Minnesota, Twin Cities
yulonglu@umn.edu

MS48

A Diffusion Model-Based Data Assimilation Framework for Uncertainty Reduction in Learning Data-Driven Dynamical Systems

We present a hybrid data assimilation framework that combines long short- term memory (LSTM) neural networks with the Ensemble Score Filter (EnSF) for nonlinear dynamical systems and operator learning. LSTM models provide accurate short-term forecasts, but their performance deteriorates over longer horizons due to error accumulation. To overcome this limitation, observational data are assimilated using EnSF, which updates the ensemble of LSTM forecasts while accommodating nonlinear and non-Gaussian state distributions. Numerical experiments with the 20-dimensional Lorenz 96 system demonstrate that the LSTMEnSF framework effectively corrects long-term forecast errors and yields robust state estimation, highlighting its promise for improving data-driven prediction in complex dynamical systems. Our results demonstrate EnSFs comparative effectiveness in the realm of traditional scientific machine learning.

Jingqiao Tang
Florida State University
jt21v@fsu.edu

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

MS48

Generative Modeling in Partially Observed Dynamical Systems

We present a measure transport approach for solving inverse problems sequentially in stochastic dynamical systems. By utilizing a Bayesian framework, we construct a generative modeling method that maps the prior processes to amortized posterior processes. This can then generate data-assimilated trajectories and perform downstream tasks such as time series prediction and system identification.

Nicole Yang
Emory University
nicoley031@gmail.com

MS50

Uncertainty-aware Virtual Sensing for Critical Energy Infrastructure with Distribution-free Guarantees

Neural operator networks provide efficient surrogates for complex physical systems, but their adoption in safety-critical domains is restricted by the lack of calibrated, spatially resolved uncertainty quantification (UQ). Accuracy alone is insufficient: an incorrect prediction with unwarranted confidence can be catastrophic. Risks grow under sparse sensing, distributional shifts, and unseen conditions, where even advanced operators may fail silently. Existing UQ methods fall short. Bayesian neural networks are costly and fragile to priors. Deep ensembles improve coverage but require heavy training and memory. Quantile regression is lightweight yet uncalibrated. What good is accuracy if predictions cannot be trusted? We introduce the Conformalized Monte Carlo Operator (CMCO), the first framework to deliver distribution-free, spatially resolved UQ for neural operators with guaranteed conditional coverage using a single model. CMCO combines Monte Carlo dropout with split conformal prediction to yield calibrated intervals without retraining or ensembling, and applies directly to DeepONet and S-DeepONet. Demonstrations on turbulent flows, elastoplastic solids, and global radiation mapping confirm robust coverage. By embedding distribution-free calibration into operator learning, CMCO makes UQ a core capability of neural surrogates for safety-critical science and engineering.

Kazuma Kobayashi
University of Illinois Urbana-Champaign
kazumak2@illinois.edu

Shailesh Garg
Department of Applied Mechanics
Indian Institute of Technology Delhi
amz218308@am.iitd.ac.in

Seid Koric, Syed Bahauddin Alam
University of Illinois Urbana-Champaign
koric@illinois.edu, alams@illinois.edu

MS50

Bayesian NNs with Meaningful Functional Priors

Bayesian neural networks (BNNs) quantify the uncertainty in the output of a neural network by constructing a Bayesian posterior distribution over the network parameters. However, the prior, which is of key importance in any Bayesian setting, is rarely meaningful for BNNs. This is because the complexity of the input-to-output map of a BNN makes it difficult to understand how certain distributions enforce any interpretable constraint on the output space. Gaussian processes (GPs), on the other hand, are often preferred in uncertainty quantification tasks due to their interpretability. The drawback is that GPs are limited to small datasets without considerable effort. To combine the two techniques, we introduce a new class of priors for BNNs which approximates a desired GP. Unlike previous approaches, our method works by defining a prior directly over the network parameters from the Mercer rep-

resentation of the covariance kernel, and does not rely on the network having a specific structure. In doing so, we can exploit the scalability of BNNs in a meaningful Bayesian way.

Alex Alberts, Ilias Bilionis
Purdue University
albert31@purdue.edu, ibilion@purdue.edu

MS50

Constrained Conditional Flow Matching Algorithms for the Solution of Probabilistic Inverse Problems

Flow matching models are generative models that use neural networks to approximate the velocity field corresponding to fixed conditional probability paths between a reference distribution and a desired target distribution. This velocity field can be shown to satisfy the appropriate continuity equation that describes the transport of the probability density functions of the reference distribution and the posterior distribution. In this work, we show that flow matching is a powerful tool for solving inverse problems within the Bayesian paradigm. Flow matching offers several advantages, such as the flexibility of choosing the reference distribution, including complex data-driven priors, accommodating intractable likelihoods, and amortizing the cost of inference. Specifically, we compare flow matching models built using the Gaussian distribution and the prior for the inverse problem as the reference distribution, where the latter represents the prior to posterior transport map. In this work, we also explore how constraints can be enforced within the framework of flow matching such that the final solution continues to be faithful to the posterior distribution for the inverse problem. Results show that flow matching can be an effective framework for solving probabilistic inverse problems with constraints.

Assad A. Oberai, Ali Fardisi, Mehrnegar Aminy
University of Southern California
aoberai@usc.edu, fardisi@usc.edu, maminy@usc.edu

MS50

Neural & Polynomial Chaos: Spectral Stochastic Methods for Operator Learning

Building surrogate models for uncertainty quantification is essential for many applications where randomness is unavoidable. The Polynomial Chaos Expansion (PCE) is widely recognized as a to-go method for constructing stochastic solutions in both intrusive and non-intrusive ways. However, to date, its application as a surrogate model has been limited to a functional regression setting. In this work, we first extend the PCE formulation to operator learning—demonstrating that PCE can provide an extremely computationally efficient alternative to neural network-based operator learning methods. However, PCE may not afford the flexibility that neural networks provide. For example, PCE requires specialized treatments to manage random variables that do not follow specified families of distributions or are not marginally independent. In this work, we then adopt the same formalism as the spectral expansion used in PCE; however, we replace the classical polynomial basis functions with neural networks to leverage their expressivity. To achieve this, we propose an algorithm that identifies orthogonal neural network-based basis functions in a purely data-driven manner, without any prior assumptions about the joint distribution of the random variables involved. This approach may offer greater flexibility for complex stochastic models. We demonstrate these schemes through several numerical examples of varying complexity.

Michael D. Shields, Himanshu Sharma
Johns Hopkins University
michael.shields@jhu.edu, hsharm10@jhu.edu

Bahador Bahmani
Northwestern University
bahador.bahmani@northwestern.edu

Lukas Novak
Brno University of Technology
lukas.novak4@vut.cz

Yannis Kevrekidis
Johns Hopkins University
yannisk@jhu.edu

MS51

Characterizing and Improving In-context Learning and Generalization in Transformer Architectures

Transformers perform in-context learning as an implicit least-squaresstyle update over prompt-induced features, enabling rapid adaptation from a few examples. Generalization is governed by taskfeature alignment, covariance spectra/effective rank, and noise. We first present rigorous error bounds and scaling relationships for the linear case. We use this intuition to address task superposition in the non-linear setting: one prompt can make the model solve multiple tasks simultaneously, creating both synergy and interference. We then characterize generalization error for this case via simple spectral diagnostics and improve performance by triggering conditional multi-step updates when they provably reduce error.

Karthik Duraisamy, Samet Oymak
University of Michigan Ann Arbor
kdur@umich.edu, oymak@umich.edu

MS51

Transformers for Hybrid Dynamics: A UQ Perspective

Transformers have recently emerged as foundation models for time series prediction, but their application to hybrid dynamicswhere discrete transitions interact with continuous flowsremains an open challenge. We propose a transformer-based framework designed with structures suited for capturing both types of dynamics. From an uncertainty quantification perspective, our study will compare how uncertainty propagates through interpolation and extrapolation in hybrid dynamics, relative to existing models. A further focus is on how incorporating HamiltonJacobiBellman (HJB) formulations with regularizations and noise injections can provide mathematically informed structure, potentially improving robustness and offering new insight into uncertainty in predictive modeling.

Hyemin Gu
University of Massachusetts Amherst
hgu@umass.edu

MS51

Icon to Genicon: Probabilistic Operator Learning

with Uncertainty Quantification

In-context operator networks (ICON) learn solution operators for ODEs/PDEs by conditioning on example initial/boundary data and their solutions. Ill present a probabilistic interpretation: ICON implicitly computes the posterior predictive mean given the context. Using random differential equations, this connects ICON to Bayesian inference and naturally leads to a generative variant, GenICON, that samples from the posterior predictive for principled uncertainty quantification. The framework unifies operator learning under a Bayesian lens and delivers uncertainty-aware predictions.

Siting Liu
UC Riverside
siting.liu1@ucr.edu

MS52

Generative Modeling without Clean Data: Self-Consistent Transport under Black-Box Corruptions

Generative modeling aims to learn an underlying data distribution from samples. In many scientific and engineering settings, however, clean samples are never observed; instead, data are available only after passing through a noisy, possibly nonlinear and ill-conditioned corruption channel. The challenge is therefore to learn a generative model for the clean distribution using only corrupted observations and access to the forward process. In this talk, I introduce the Self-Consistent Stochastic Interpolant (SCSI), a transport-based framework that inverts such corruption channels at the level of distributions. The method iteratively refines a transport map so that, when composed with the forward model, it reproduces the observed corrupted distribution. This fixed-point formulation yields an efficient and flexible algorithm requiring only black-box evaluations of the forward operator. We establish convergence guarantees under suitable assumptions and demonstrate strong empirical performance on high-dimensional problems in imaging and scientific reconstruction. Joint work with Chirag Modi, Eric Vanden-Eijnden, and Joan Bruna (arXiv:2512.10857).

Jiequn Han
Flatiron Institute
jhan@flatironinstitute.org

MS52

Diffeomorphism Invariant Measure Flows in Bayes Hilbert Spaces

Measure flowstime-parametrized evolutions of measuresare now central to generative modeling, sampling, and Bayesian inverse problems; however, many practically useful dynamics fall outside the classical metric gradient-flow formalism. In this work, we develop an alternative framework based on Bayes Hilbert spaces. We first extend the theory of Bayes Hilbert spaces by giving a new construction using a Gelfand triple which brings singular Radon measures into the theory while retaining a compatible Hilbert structure. We apply this framework to study diffeomorphism invariant flows, which enjoy convergence rates that are independent of the target measure. We prove two structural results: (i) pushforward by a diffeomorphism is an isometric isomorphism between Bayes Hilbert spaces, meaning the geometry of Bayes Hilbert spaces itself is diffeomorphism invariant, and (ii) the FisherRao gradient flow of KullbackLeibler divergence, the only diffeomorphism invariant flow in the classical metric gradient flow framework, coincides with straight-line interpolation in a Bayes Hilbert space. Building on these, we introduce new diffeomorphism-invariant dynamicsq-flows and Bayes Hilbert splinesthat are more flexible than the Fisher-Rao flow but still preserve full diffeomorphism invariance. Numerical simulations for sampling and inverse problems demonstrate the practical utility of these new flows.

David Mis, Maarten V. de Hoop
Rice University
david.mis@rice.edu, mdehoop@rice.edu

MS52

GenFocal: Generative Downscaling for Climate Risk Assessment and Data Assimilation

Assessing climate risk is critical to decision-making across many social and economic domains, but it requires accurate, high-resolution information that coarse global climate models, the workhorse of climate research, cannot provide. Bridging this resolution gap in a computationally efficient way is a major scientific challenge. In this talk we introduce GenFocal, an AI framework coupling flow-matching and diffusion models to generate statistically accurate, fine-scale weather from coarse climate projections, without requiring paired simulated and observed events during training. GenFocal synthesizes complex and long-lived hazards, such as heat waves and tropical cyclones, even when they are not well represented in the coarse climate projections. It also samples high-impact, rare events more accurately than leading methods. We will introduce the framework, including its rationale, and we will show several examples showcasing its performance, which hinges on translating large-scale climate projections into actionable, localized information that can be used to assess regional climate risks.

Leonardo Zepeda-Núñez
Google
lzepedanunez@google.com

MS52

On the Use of Diffusion Priors for Scientific Inverse Problems

This talk focuses on algorithms that utilize pretrained diffusion models as priors for inverse problems. The first section introduces InverseBench, a study of various plug-and-play diffusion prior algorithms developed recently. We assess and compare these algorithms across five distinct scientific inverse problems, offering new insights into their effectiveness and limitations. The second section addresses the often-overlooked aspect of uncertainty calibration. We present Blade, a derivative-free method for calibrated Bayesian inversion. We establish convergence analysis and provide explicit error bounds to quantify how approximations in the learned score function and statistical linearization affect the sampling. On both controlled linear setting and challenging nonlinear fluid dynamics problems, we show that Blade is capable of producing well-calibrated posterior samples while existing derivative-free methods fail.

Hongkai Zheng
California Institute of Technology

hzzheng@caltech.edu

MS53

Structure-Informed Uncertainty Quantification in Generative Models

I will present a novel framework for learning unknown data distributions with known underlying structures. I will rigorously demonstrate how incorporating the structure into models can reduce the model uncertainty along with a new Wasserstein Uncertainty Propagation result on unbounded domains with sample complexity. This framework applies to both SDE-based generative models and deterministic flow models, and may fit in broader scenarios with physical constraints. As a particular example, I will show that equivariant diffusion models can provably learn distributions with symmetries efficiently. This is a joint work with Markos Katsoulakis and Benjamin Zhang.

Ziyu Chen
University of North Carolina at Chapel Hill
School of Data Science and Society/Department of Mathematics
ziyuchen@unc.edu

MS53

Physics-Informed Laplace Neural Operator for Learning Differential Equations

Differential equations are fundamental for modeling complex scientific and engineering processes. Operator learning offers a powerful framework for approximating the mapping from parametric inputs such as initial/boundary conditions or forcings to the corresponding solutions. However, purely data-driven operator learning methods such as the Laplace Neural Operator (LNO) typically require large training datasets and exhibit limited extrapolation capabilities. We propose the Physics-Informed Laplace Neural Operator (PILNO), which embeds the governing physical laws directly within the learning process. This physics-informed approach reduces data dependency, enabling effective learning insmall-data regimesand enhancinggeneralization to out-of-distribution inputswhere purely data-driven LNOs often fail. By leveraging both data (when available) and physics, PILNO offers a more robust and data-efficient pathway to learning solution operators. We demonstrate advantages of PILNO across diverse parametric differential equation problems, highlighting its improved data efficiency and extrapolation performance.

Minseok Choi
POSTECH Mathematics Institute, Korea
mchoi@postech.ac.kr

MS53

Stable Derivative Free Gaussian Mixture Variational Inference

This talk is concerned with the approximation of probability distributions known up to normalization constants, with a focus on Bayesian inference for large-scale inverse problems in scientific computing. In this context, key challenges include costly repeated evaluations of forward models, multimodality, and inaccessible gradients for the forward model. To address them, we develop a variational inference framework that combines Fisher-Rao natural gradient with specialized quadrature rules to enable derivative free updates of Gaussian mixture variational families. The resulting method, termed Derivative Free Gaussian Mixture Variational Inference (DF-GMVI), guarantees covariance positivity and affine invariance, offering a stable and efficient framework for approximating complex posterior distributions. The effectiveness of DF-GMVI is demonstrated through numerical experiments on challenging scenarios, including distributions with multiple modes, infinitely many modes, and curved modes in spaces with up to 100 dimensions. The method's practicality is further demonstrated in a large-scale application, where it successfully recovers the initial conditions of the Navier-Stokes equations from solution data at positive times.

Zhengyu Huang
Peking University
huangdz@bicmr.pku.edu.cn

MS53

A Multi-Fidelity Method for Velocity Discretization of Boltzmann Equations

In this talk, we introduce a bi-fidelity algorithm for velocity discretization of Boltzmann-type kinetic equations under multiple scalings. The proposed method employs a simpler and computationally cheaper low-fidelity model to capture a small set of significant velocity points through the greedy approach, then evaluates the high-fidelity model only at these few velocity points and to reconstruct a bi-fidelity surrogate. This method, inspired by its previous application in solving uncertainty quantification problem for kinetic models, integrates a simpler collision term of relaxation type in the low-fidelity model and an asymptotic-preserving scheme in the high-fidelity update step. Both linear Boltzmann under diffusive scaling and the nonlinear full Boltzmann in hyperbolic scaling are discussed. The weak asymptotic-preserving property and empirical error bound estimates are shown. Extensive numerical experiments on Boltzmann problems with smooth or discontinuous initial conditions and under various regimes have been carefully studied, which demonstrates the effectiveness and efficiency of our designed numerical schemes.

Liu Liu
Chinese University of Hong Kong
liuliu@cuhk.edu.hk

MS54

Fast and Accurate Spectral Methods for Irregularly Sampled Data

A common problem in exploratory data analysis is wanting to use Fourier analysis and the FFT, but being unable to due to gaps or irregularities in sampling. Many intuitive options, like smoothing to a grid and transforming or simple window-based weighted Fourier sums, have very poor properties. In this work, I will introduce a new estimator that in many cases works as well as Fourier analysis of gridded data with very few caveats: the resulting estimator can be computed in quasilinear time, analysis of finite-sample bias is natural and direct, and the approach is equally applicable in multiple dimensions. After discussing the conceptual shift used in designing this estimator and the computational apparatus to rapidly and robustly compute it, I will close with some dramatic demonstrations where this method works well and alternative irregular data approaches completely fail to give correct or informative results.

Christopher Geoga

Department of Statistics
University of Wisconsin-Madison
geoga@wisc.edu

Paul G. Beckman
Courant Institute, New York University
paul.beckman@cims.nyu.edu

MS54

Convergence Analysis of a Greedy Algorithm for Conditioning Gaussian Random Variables

In the context of Gaussian conditioning, greedy algorithms iteratively select the most informative measurements, given an observed Gaussian random variable. However, the convergence analysis for conditioning Gaussian random variables remains an open problem. We adress this by introducing an operator M that allows us to transfer convergence rates of the observed Gaussian random variable approximation onto the conditional Gaussian random variable. Furthermore we apply greedy methods from approximation theory to obtain convergence rates. These greedy methods have already demonstrated optimal convergence rates within the setting of kernel based function approximation. In this paper, we establish an upper bound on the convergence rates concerning the norm of the approximation error of the conditional covariance operator.

Daniel Winkle
University of Stuttgart
daniel.winkle@mathematik.uni-stuttgart.de

MS54

Bilevel Optimization for Learning Hyperparameters: Application to Solving Pdes and Inverse Problems with Gaussian Processes

The efficacy (e.g., accuracy, stability, generalization) of most methods for solving scientific computing and inference problems strongly depends on the choice of hyperparameters. Such methods include kernel- and neural network-based approaches for partial differential equations (PDEs), inverse problems, and supervised learning tasks. While bilevel optimization offers a principled framework for hyperparameter tuning, its nested optimization structure can be computationally demanding, especially in PDE-constrained contexts. In this paper, we propose an efficient strategy for hyperparameter optimization within the bilevel framework by employing a Gauss-Newton linearization of the inner optimization step. Our approach provides closed-form updates, eliminating the need for repeated costly PDE solves. As a result, each iteration of the outer loop reduces to a single linearized PDE solve, followed by explicit gradient-based hyperparameter updates. We demonstrate the effectiveness of the proposed method through Gaussian process models applied to nonlinear PDEs, PDE systems, and inverse problems. Extensive numerical experiments highlight substantial improvements in accuracy and robustness compared to conventional random hyperparameter initialization. In particular, experiments with additive kernels and neural network-parameterized deep kernels highlight the method's scalability and effectiveness for high-dimensional hyperparameter optimization.

Xianjin Yang
California Institute of Technology
yxjmath@caltech.edu

MS55

Information Field Theory Approach to State Estimation and Calibration of Stochastic Differential Equations

Current smoothing methods for stochastic differential equations (SDEs) scale poorly with the curse of dimensionality. In this work, we introduce an information field theory (IFT) approach for smoothing SDEs. We identify a formal prior probability measure in function space which describes the law of the SDE. Under IFT this prior can be conditioned on observations to solve the smoothing problem. To numerically characterize the associated posterior, we restrict the prior to a finite-dimensional function space in the form of basis functions. In doing so, our method scales with the number of observations, rather than the number of time steps, mitigating the curse of dimensionality.

Maxwell Bolt, Alex Alberts
Purdue University
boltm@purdue.edu, albert31@purdue.edu

Ilias Bilionis
School of Mechanical Engineering
Purdue University
ibilion@purdue.edu

MS55

The Ensemble Kalman Inversion Race

Iterative ensemble methods were initially created to solve data assimilation in the geosciences but are now popular in applications beyond their original use cases. One application of particular interest is calibrating climate model parameters (e.g., cloud microphysics parameterizations) to climate statistics. Many variations of iterative ensemble methods exist, including variants of ensemble Kalman inversion and the ensemble smoother. Given the high computational cost associated with learning climate model parameters and assuming that the various ensemble methods have different relative performances, an important question one should ask is: Which iterative ensemble method should I use for climate model calibration? We consider several ensemble methods and compare them on a series of systematic numerical experiments where we estimate model parameters from statistical observations. All test problems are characterized by noisy loss functions, with a noise level too high for successful gradient-based optimization. Our results indicate unscented methods are best for problems with low-dimensional parameter spaces. In high-dimensional parameter spaces, the iterative ensemble Kalman filter appears to be most efficient when informative prior information is also available. With uninformative prior information, ensemble transform methods are the most robust and efficient, making ensemble transform methods promising for learning climate model parameters from statistics.

Rebecca Gjini
Scripps Institution of Oceanography
University of California, San Diego
rgjini@ucsd.edu

Matthias Morzfeld
University of California, San Diego
Scripps Institution of Oceanography

mmorzfeld@ucsd.edu

Oliver Dunbar, Tapio Schneider
California Institute of Technology
odunbar@caltech.edu, tapio@caltech.edu

MS55
Multistage Neural Network for Inverse Problems

Deep learning techniques are increasingly being applied to scientific challenges, particularly inverse problems, where network precision is crucial. In practice, however, existing scientific machine learning approaches often struggle to model complex, multi-scale dynamical systems with the high precision required in scientific contexts. To address this issue, we developed Multi-Stage Neural Networks (MSNN), a method that divides the training process into sequential stages. Each stage introduces a new network optimized to fit the residual error from the previous one. In this talk, I will use illustrative examples to explain the principles and features of our method for both standard regression tasks and physics-informed neural networks. Finally, I will present a concrete example to demonstrate the potential of this approach for solving challenging inverse problems.

Yongji Wang
The Courant Institute of Mathematical Sciences
yw8211@nyu.edu

Ching-Yao Lai
Stanford University
cyaolai@stanford.edu

MS56
Explaining Models under Multivariate Bernoulli Distribution via Hoeffding Decomposition

The analysis and explainability study of predictive models with random inputs can be achieved by decomposition sub-models having easier interpretable features. Arising from the uncertainty quantification community, recent results have demonstrated the existence and uniqueness of a generalized Hoeffding decomposition for such predictive models when the stochastic input variables are correlated, based on concepts of oblique projection onto $\mathbb{L}^2$ subspaces. This work focuses on the case where the input variables have Bernoulli distributions and provides a complete tool for description of this decomposition. We show that in this case the underlying $\mathbb{L}^2$ subspaces are one-dimensional and that the functional decomposition is explicit. This leads to a complete interpretability and allows reverse engineering. Numerous explicit indicators that help to understand the influence of inputs on the output prediction (such as Sobol' indices, Shapley effects, etc.) are explicitely derived. Illustrated by numerical experiments, this type of analysis proves useful for addressing decision-support problems, based on binary decision diagrams, Boolean networks or binary neural networks. Beyond the case of binary inputs, the paper outlines perspectives for extending these findings to models with finite countable inputs.

Baptiste Ferrere
EDF R&D
baptiste.ferrere@edf.fr

Nicolas Bousquet
EDF & Sorbonne Université
nicolas.bousquet@edf.fr

Fabrice Gamboa
Institut de Mathématiques de Toulouse. AOC Project
University of Toulouse
fabrice.gamboa@math.univ-toulouse.fr

Jean-Michel Loubes
Equipe de Statistique et Probabilités UMR 5219
Université de Toulouse 3, Institut de Mathématiques
loubes@math.univ-toulouse.fr

Joseph Muré
EDF R&D
joseph.mure@edf.fr

MS56
Gradient-Enhanced Global Sensitivity Analysis with Poincaré Chaos Expansions

Chaos expansions are widely used in global sensitivity analysis (GSA), as they leverage orthogonal bases of L^2 spaces to efficiently compute Sobol indices, particularly in data-scarce settings. When derivatives are available, we argue that a desirable property is for the derivatives of the basis functions to also form an orthogonal basis. We demonstrate that the only basis satisfying this property is the one associated with weighted Poincaré inequalities and Sturm-Liouville eigenvalue problems, which we refer to as the Poincaré basis. We then introduce a comprehensive framework for gradient-enhanced GSA that integrates recent advances in sparse, gradient-enhanced regression for surrogate modeling with the construction of weighting schemes for derivative-based sensitivity analysis. The proposed methodology is applicable to a broad class of probability measures and supports various choices of weights. We illustrate the effectiveness of the approach on a challenging flood modeling case study, where Sobol indices are accurately estimated using limited data.

Olivier Roustant
INSA Toulouse
roustant@insa-toulouse.fr

Nora Luethen
ETH Zurich
luethen@ibk.baug.ethz.ch

David Heredia
INSA Toulouse
dheredia@insa-toulouse.fr

Bruno Sudret
ETH Zurich
Chair of Risk, Safety and Uncertainty Quantification
sudret@ibk.baug.ethz.ch

MS56
The Csiszr Index and Variable Ranking

We first introduce the class of f-divergences between two probability measures along with several classical examples. We then define the Csiszr index for a pair of random variables (X, Y) as the f-divergence between the product of their marginals and their joint distribution. This index can be seen as a sensitivity index based on a measure of the dependence. We show that the f-divergence is invariant under a general class of transformations. From this result, we derive a condition for the invariance of the Csiszrs index. For random vectors (X, Y) with atoms, the associated

copula is no longer unique. We establish that the Csiszrs index can be reformulated as an f-divergence between the checkerboard product copula and the checkerboard joint copula. More generally, for any copula of (X, Y), the Csiszrs index is greater than or equal to the f-divergence between the product copula and the joint copula. We develop the case where (X, Y) is a Gaussian vector. When both X and Y are univariate, the Csiszr index is a non decreasing function of the absolute value of their correlation, for any suitable function f. This result is then extended to higher-dimensional settings and applied to variable ranking.

Michael Baudin
EDF R&D
michael.baudin@edf.fr

Cristina Butucea
ENSAE
cristina.butucea@ensae.fr

Jean-Francois Delmas
Ecole des Ponts ParisTech
delmas@cermics.enpc.fr

Anne Dutfoy, Antoine Schoonaert
EDF R&D
Anne.Dutfoy@edf.fr, antoine.schoonaert@edf.fr

MS57

Active Learning Methods for Multi-Output Uncertainty Quantification Based on Gaussian Process Regression and Principal Component Analysis

In recent years, surrogate models have been increasingly proposed to speed-up uncertainty propagation tasks in many domains of science and technology. Among the various methods, Gaussian process regression (GPR) stands as a premier tool, providing accurate models and an indication of their prediction confidence. The confidence information also supports the iterative acquisition of an optimal sequence of training points in so-called active learning (AL) schemes. This aspect is of paramount importance in complex engineering problems, since the computational cost incurred by a single simulation often requires the number of training samples to be minimized. However, standard AL methods mostly address deterministic or optimization scenarios, while becoming non-optimal when input parameters are equipped with a probability distribution. In this work, we introduce an active learning (AL) framework specifically designed for uncertainty propagation, in which the distribution of the input parameters is incorporated in the definition of the acquisition function that selects the optimal sampling points. Moreover, to deal with multi-output problems, the training dataset is compressed using principal component analysis (PCA) and the acquisition criterion is applied to the least-accurate principal component. The method is applied to partial differential equations with stochastic coefficients, as well as to the frequency-domain simulation of an electromagnetic structure.

Michele Cusano
Department of Electronics and Telecommunications
Politecnico di Torino
michele.cusano@polito.it

Paolo Manfredi
Politecnico di Torino
paolo.manfredi@polito.it

MS57

Data-based Low-order Modelling for Optimal Control under Uncertainties

We consider stochastic dynamics with metastable states as they serve as models in many scientific disciplines like in simulations of macro-molecules, in climate systems, and in applications of uncertainty quantification. Here, metastability refers to a fixed number relatively pertinent macro-states (as compared to the permanent stochastic fluctuations) in a dynamical systems state space. In particular, the transitions between these macro-states are rare events. In this work, we propose a model of reduced complexity based on the so-called extended dynamic mode decomposition (EDMD) and on biased sampling. As a result, we obtain a bilinear in control and state model that solely derives from data and that can be employed as a surrogate model for optimal control or uncertainty quantification. We present the underlying ideas and show numerical examples that illustrate the performance and functionality.

Jan Heiland
Institute of Mathematics
Technical University Ilmenau
jan.heiland@tu-ilmenau.de

Lei Guo
lguo@mpi-magdeburg.mpg.de
mpi for dynamics of complex technical systems

Fan Wang
MPI for Dynamics of Complex Technical Systems
fwang@mpi-magdeburg.mpg.de

Feliks Nüske
Max Planck Institute for Dynamics of
Complex Technical Systems, Magdeburg, Germany
nueske@mpi-magdeburg.mpg.de

MS57

Methods Based on Polynomial Chaos for Delay Differential Equations Including Random Variables

We consider initial value problems of systems consisting of delay differential equations (DDEs). Physical parameters of a system are replaced by random variables to conduct an uncertainty quantification. Consequently, the solution of a DDE system becomes a random process, which has to be computed numerically. We expand the random process into a series of the generalized polynomial chaos, where predetermined basis polynomials as well as time-dependent coefficient functions are included. Now the aim is to compute a finite set of the unknown coefficient functions. On the one hand, stochastic collocation techniques produce approximations by solving the DDEs for a finite number of realizations of the random variables. On the other hand, a stochastic Galerkin method yields a larger deterministic system of DDEs, which has to be solved once to obtain an approximation of the coefficient functions. We investigate the application of both approaches to solve the stochastic model. If the stochastic Galerkin method is structure-preserving, then the larger system of DDEs can be solved by the same numerical methods as the original DDEs. We present results of numerical computations using an illus-

trative example.

Roland Pulch
University of Greifswald
pulchr@uni-greifswald.de

MS57

Predictive Reduced-Order Modeling in Coupled Systems of PDEs

Fluid structure interaction problems can be highly sensitive to the values of physical parameters. For instance, solvers may experience instabilities due to the added mass effect when fluid and structure densities are similar, and a poroelastic structure modeled by Biots equations can exhibit pressure oscillations or Poisson locking with small permeability constants or large values of the Lam parameter λ. We present a strongly coupled partitioned method for the solution of coupled PDEs based on a well-posed monolithic formulation of the system which employs one or more Lagrange multipliers (LMs) representing interfacial flux quantities. The partitioned scheme is centered around a Schur complement equation, whose solution at each time step allows for the decoupling of the respective subproblems, making the method non-iterative between subdomains. To increase computational efficiency, projection-based reduced order models (ROMs) are introduced into one or both subdomains. We examine the methods performance with both linear elastic and poroelastic structures over a range of reproductive and predictive parametric tests, demonstrating the robustness of the coupling scheme as well as the viability of ROMs to improve computational performance.

Amy de Castro
Department of Mathematics
University of Utah
u6067603@utah.edu

MS58

Low Discrepancy Drop-In Replacements for IID

Low discrepancy (LD) sequences facilitate faster convergence for computing expectations or integrals of high dimensional functions. Unfortunately, the best low discrepancy sequences have preferred sample sizes, n, such as integer powers of 2. This may be inconvenient for practitioners who are accustomed to using independent and identically distributed (IID) samples that have no restrictions on n. This talk describes recent work on lattice, Kronecker, and other LD sequences where the n may be any positive integer, and the dimension, d, is also allowed to be of arbitrary size. We use component-by-component constructions based on a figure of merit that includes the discrepancies for all values of n up to a maximum value, N. Because these discrepancies are based on carefully chosen kernels, the cost of the constructions is at most $\mathcal{O}(dN\log(N))$.

Fred J. Hickernell
Illinois Institute of Technology
Dept of Applied Math & Ctr for Interdisc Sci Comput
hickernell@illinoistech.edu

Claude Hall, Jr.
Illinois Institute of Technology
Chicago, IL
chall14@illinoistech.edu

Larysa Matiukha
Illinois Institute of Technology
Department of Applied Mathematics
lmatiukha@illinoistech.edu

Jimmy Nguyen
University of California, Irvine
jimmyhn7@uci.edu

Alexander Pride
Illinois Institute of Technology
Department of Applied Mathematics
apride@hawk.illinoistech.edu

MS58

ESPs: Cost-Efficient Sampling of Computationally Expensive Posterior Distributions

Computationally expensive posterior distributions arise in modern scientific settings. One such setting is Bayesian inverse problems with computer simulators, where each evaluation of the (unnormalized) posterior p requires a forward run of the expensive simulator that can take hundreds of CPU hours. This cost poses a challenge for existing posterior samplers, where each sample requires at least one evaluation of the posterior p. Given a computational budget, a "cost-efficient" sampler is desired for effective posterior exploration with limited posterior evaluations. We propose a new sampling method called cost-Efficient Stein Points (ESPs). ESPs extend the Stein points in Chen et al. (2018), which employ the sequential minimization of the kernel Stein discrepancy to generate a sequence of posterior samples $x_1, x_2, \cdots$. To reduce posterior evaluations, we leverage a Gaussian process surrogate on the kernel discrepancy that guides the selection of evaluation points. This facilitates the optimization of samples $x_1, x_2, \cdots$ in a cost-efficient manner, where all previous posterior evaluations can be re-used for optimizing the current sample. We prove that, under mild conditions, ESPs converge in distribution to the desired posterior distribution. Finally, we demonstrate the cost-efficiency of ESPs over the state-of-the-art in numerical experiments and an application for Bayesian calibration of a biological oscillator process.

Simon Mak
Duke University
sm769@duke.edu

MS58

Multi-Fidelity Quasi-Monte Carlo Sampling

Multifidelity Monte Carlo (MFMC) methods provide an effective framework for reducing the computational cost of uncertainty quantification by combining information from inexpensive low-fidelity models with corrections from more accurate high-fidelity simulations. Traditionally, MFMC relies on standard Monte Carlo sampling, and its efficiency is therefore limited by the slow convergence rate of Monte Carlo estimators. It is well known that quasi-Monte Carlo (QMC) techniques can significantly accelerate convergence when applied to smooth integrands. In this work, we extend the multilevel quasi-Monte Carlo methodology to the multi-fidelity setting, while avoiding costly stochastic replications for error estimation. Numerical experiments demonstrate that the proposed multi-fidelity QMC approach can achieve substantial gains in accuracy per unit cost compared to both standard MFMC and classical QMC.

Pieterjan Robbe

Sandia National Laboratories
pmrobbe@sandia.gov

Aleksei Sorokin
Illinois Institute of Technology
asorokin@hawk.iit.edu

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS59

KOH-GPJax: Bayesian Calibration of Computer Models with GPU Acceleration

Kennedy O'Hagan's (KOH) seminal paper on Bayesian Calibration of Computer Models is 25 years old and is still considered the gold standard for performing parameter tuning on physics-based simulation models. Introducing KOH-GPJax, a modern python implementation of the KOH model. We employ JAX's scientific computing prowess enabling just-in-time compilation, automatic differentiation and GPU/TPU acceleration of the KOH model. This has enabled the calibration of multiple parameters, fitting to larger datasets and integrating a variety of emulator methods from the GPJax framework. KOH-GPJax is also agnostic to the model fitting procedure allowing users to experiment with their favourite MCMC algorithm or package. This package is enabling a new class of Bayesian calibration problems has a growing community of users including researchers at the Universities of Cambridge, Stanford and Texas.

James Briant, Matthew Graham, Serge Guillas
University College London
james@briant.co.uk, m.graham@ucl.ac.uk, s.guillas@ucl.ac.uk

MS59

Robustgasp, RobustCalibration: Two Packages for Forward Prediction and Inverse Estimation

In this talk, we introduce RobustGaSP and RobustCalibration packages for forward prediction and inversion estimation. The first package implements the parallel partial Gaussian process surrogate model for emulating expensive computer simulations with either scalable output or massive coordinates. The tool is implemented in the RobustGaSP package available in R, MATLAB, and Python, for predicting both scalar- and vector-valued outputs with uncertainty assessment through the predictive standard deviation. The second tool is implemented in the RobustCalibration package, which handles Bayesian data inversion or model calibration by one or multiple types of experimental observations. A unique feature of the package is the inclusion of fast surrogate models of both scalar- and vector-valued computer simulations that bypass the expensive simulation in one line of code. Additionally, any difference between the simulation and reality can be incorporated by modeling the discrepancy function as a Gaussian process or scaled Gaussian process model. These tools have been applied in a wide range of science and engineering research problems that include emulating expensive simulation at molecular-, meso- and macro-scales, active learning with error control, and data inversion with satellite interferograms or velocity fields from microscopy data.

Clayton Ellis
University of California, Santa Barbara
claytonellis@ucsb.edu

MS60

Goal-Oriented OED for Bayesian Inverse Problems Governed by Autonomous Dynamical Systems

We present a goal-oriented algorithm for optimal experimental design (OED) that extends our recent work on real-time, extreme-scale Bayesian inversion for linear autonomous dynamical systems. That work utilized a predetermined set of observation points (sensors) to perform real-time parameter inference and predict quantities of interest (QoIs). In practice, sensor deployment is often constrained by a budget, making the strategic placement of a limited number of sensors a critical design challenge. We address this challenge with a framework for optimal sensor subset selection. Our approach employs a greedy algorithm that, given a set of candidate locations, iteratively selects sensors to maximize a measure of expected information gain. This information gain objective is tailored to be goal-oriented, prioritizing sensor locations that most reduce uncertainty in specific QoIs. The efficiency of this framework is enabled by the fast, FFT-based operators from our prior work, and the greedy selection process is readily parallelizable on multi-GPU clusters. The goal-oriented OED framework is applied to a digital twin (DT) for tsunami early warning in the Cascadia subduction zone. This DT assimilates observations from seafloor pressure sensors during an earthquake to inform a billion-parameter inverse problem for the seafloor motion and issue uncertainty-equipped predictions of tsunami wave heights in real time. Sensor locations are optimized to enable accurate tsunami forecasts.

Stefan Henneking
University of Texas at Austin
stefan@oden.utexas.edu

Sreeram R. Venkat
University of Texas Austin
201 E 24th St, Austin, TX 78712
srvenkat@utexas.edu

Omar Ghattas
University of Texas at Austin
omar@oden.utexas.edu

MS60

A Generic Digital Twin-Enhanced Framework for Partial Discharge Diagnosis

Accurate prognostics of complex systems has become an imperative research area in the field of reliability engineering. In particular, an optimal sensor network can enhance monitoring of high-risk systems across many fields. Our study focuses on partial discharge detection and localization, which is a severe fault in many high power applications. Specifically, we aim to develop a data fusion framework enhanced with a digital twin, which can integrate data from numerous sources to achieve optimal system analysis. The essentials of this study lie within the scope of data selection for machine learning and data fusion models. However, the challenge of efficient data selection has

yet to be formulated in the context of optimal design and sensor placement. Specifically, the proposed method can simultaneously determine the sensor network while optimizing the predictor or monitoring system. This is accomplished by using proxy models for data selection, and explainable AI for model transparency. The global and local explainability are provided via SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) analyses, respectively. Various search algorithms are implemented to solve the generally applicable mathematical formulation, which utilizes predictor accuracy and interpretability as the fitness function. Overall, this proposed method is demonstrated on partial discharge detection and localization within electric motors.

Sara Kohtz
Binghamton University
skohtz1@binghamton.edu

MS60

Towards Digital Twin for Design of Experiments with Moving Sensor

We optimize the path of a mobile sensor to minimize the posterior uncertainty of a Bayesian inverse problem. Along its path, the sensor continuously takes measurements of the state, which is a physical quantity modeled as the solution of a partial differential equation (PDE) with uncertain parameters. We derive the closed-form expression of the posterior covariance matrix of the model parameters as a function of the path, and formulate the optimal experimental design problem for minimizing the posterior's uncertainty. We discretize the problem such that the cost function remains consistent under temporal refinement. Our formulation allows the addition of constraints that model the dynamics of the sensor platform and ensure that the path avoids obstacles and remains physically interpretable through a control parameterization. We present computational results for a convection-diffusion equation with unknown initial condition.

Sven Leyffer
Argonne National Laboratory
leyffer@anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@anl.gov

Nicole Aretz
The University of Texas at Austin
nicole.aretz@austin.utexas.edu

Thomas F. Lynn
Argonne National Laboratory
tlynn@anl.gov

MS60

Optimal Experimental Design for Biomanufacturing Digital Twin Calibration and Process Control

Biomanufacturing processes involve intricate interactions among biological, chemical, and physical factors across multiple scales. Designing optimal control strategies for such systems is challenging due to their nonlinear, stochastic nature, and limited data availability. To address this, we introduceActor-Simulator, a novel framework that integrates digital twin calibration into model-based reinforcement learning (MBRL) for optimal control of stochastic systems with complex dynamics and delay effects. Unlike traditional model-based approaches that rely on restrictive assumptions and overlook parameter uncertainty, Actor-Simulator jointly calibrates a multi-scale mechanistic modelformulated as stochastic differential equations (SDEs)and learns an optimal control policy. This dual-component strategy balances exploration and exploitation by using policy performance to guide data collection, targeting the most informative regions of the state-action space. To deepen mechanistic understanding, we further quantify uncertainty propagation in stochastic reaction networks using gradient-based analysis of their SDE representations. This enables rigorous sensitivity analysis of multivariate output trajectories with respect to perturbations in model parameters and decision inputs, while addressing SDE discretization bias. Our method provably converges to the optimal policy and demonstrates superior performance in biopharmaceutical manufacturing simulations.

Wei Xie
Northeastern University
w.xie@northeastern.edu

MS61

Sparsity-Promoting Priors for 3D Radar Imaging

Three-dimensional radar imaging presents a large-scale, complex-valued inverse problem, where even applying the forward operator is computationally demanding. Incorporating a third spatial dimension has the potential to dramatically improve automated target recognition (ATR), which often employs shape-based object detection algorithms. This improvement would stem from the fact that height information provides discriminative features that are difficult to capture in 2D projections. To leverage this potential, 3D reconstructions must preserve the structural integrity of objects while suppressing spurious scatterers. We investigate sparsity-promoting priors and regularization strategies that encourage faithful recovery of target shapes, including penalties in both image and transform domains. In addition, we explore techniques that enforce connectedness and neighborhood consistency after thresholding, aiming to improve interpretability and robustness.

Victor Churchill
Trinity College
victor.churchill@trincoll.edu

MS61

Signal Recovery Using a Complex-Valued Bayesian Lasso

Recovering complex-valued images from noisy indirect data is important in applications such as ultrasound and synthetic aperture radar imaging. While there are many effective algorithms to recover point estimates of the magnitude, fewer are designed to recover the phase. Quantifying uncertainty in the estimate can also provide valuable information for real-time decision making. This investigation therefore proposes a new Bayesian inference method that recovers point estimates while also quantifying the uncertainty for complex-valued signals or images given noisy and indirect observation data. Our method is motivated by the Bayesian LASSO approach for real-valued sparse signals, and we demonstrate that the Bayesian LASSO can be effectively adapted to recover complex-valued images whose magnitude is sparse in some (e.g., the gradient) domain. We present numerical examples to demonstrate our

algorithms robustness to noise as well as its computational efficiency.

Dylan Green
Trinity College
dylan.green@trincoll.edu

Jonathan Lindbloom, Anne Gelb
Dartmouth College
jonathan.t.lindbloom.gr@dartmouth.edu, annegelb@math.dartmouth.edu

MS61

Making the Most of Every Matvec: Projection-Based Methods for Efficient UQ in Sparsity-Promoting Inverse Problems

Efficient and reliable sample-based UQ for imaging inverse problems remains a challenging task. Sampling algorithms require the solution of a sequence of high-dimensional least squares problems with changing matrices and changing righthand sides. The usage of a standard algorithm such as CGLS/LSQR to solve each problem in the sequence may be prohibitively expensive due to the large number of inner iterations required, which translates to many matrix-vector products (matvecs) with the (possibly costly) forward model A and/or the sparsifying transformation L. Here, we propose a new paradigm for the efficient solution of the sequence of least squares problems based on dimensionality reduction via subspace projection. We demonstrate the efficacy of our approach when applied to several sparsity-promoting imaging inverse problems involving the LASSO (Laplace) prior, sparse bayesian learning priors, and the horseshoe prior.

Jonathan Lindbloom
Dartmouth College
jonathan.t.lindbloom.gr@dartmouth.edu

MS62

Multifidelity Bayesian Inference for Personalized Cardiac Electrophysiology Models

Personalized digital twins of the heartpatient-specific computational models of heart functionoffer a powerful framework for predictive medicine and treatment planning. However, clinical data of cardiac electrophysiology are often sparse and noisy, leading to substantial uncertainty in model calibration. At the same time, detailed electrophysiology simulations are computationally demanding and involve high-dimensional parameter spaces, which makes traditional Bayesian inference methods such as Markov chain Monte Carlo (MCMC) infeasible. To address these challenges, we investigate a multifidelity Bayesian inference approach that combines surrogate models with detailed simulations inside a MetropolisHastings algorithm. The goal is to reduce the computational burden of posterior sampling while maintaining sufficient accuracy for uncertainty quantification. The talk will focus on the formulation of this methodology and its application to synthetic data as a testbed for scalable Bayesian inference in cardiac digital twin models.

Marie Cloet
Department of Computer Science
KU Leuven
marie.cloet@kuleuven.be

Maarten Volkaerts
Katholieke Universiteit Leuven
maarten.volkaerts@kuleuven.be

Giovanni Samaey
KU Leuven
giovanni.samaey@kuleuven.be

MS62

Rapid Calibration of Atrial Electrophysiology Models Using Gaussian Process Emulators in the Ensemble Kalman Filter

Atrial fibrillation (AF) is a common cardiac arrhythmia characterised by disordered electrical activity in the atria. The standard treatment is catheter ablation, which is invasive and irreversible. Recent advances in computational electrophysiology offer the potential for patient-specific models, often referred to as digital twins, that can be used to guide clinical decisions. To be of practical value, we must be able to rapidly calibrate physics-based models using routine clinical measurements. We pose this calibration task as a static inverse problem, where the goal is to infer tissue-level electrophysiological parameters from the available observations. To make this tractable, we replace the expensive forward model with Gaussian process emulators (GPEs) and propose a novel adaptation of the ensemble Kalman filter (EnKF) for static non-linear inverse problems. The approach yields parameter samples that can be interpreted as coming from the best Gaussian approximation of the posterior distribution. We compare our results with those obtained using Markov chain Monte Carlo (MCMC) sampling and demonstrate the potential of the approach to enable near-real-time patient-specific calibration, a key step towards predicting outcomes of AF treatment within clinical timescales. The approach is readily applicable to a wide range of static inverse problems in science and engineering.

Mariya Mamajiwala
University of Nottingham
mariya.mamajiwala@nottingham.ac.uk

MS62

Robust and Efficient Uncertainty Quantification and Software Tools for Digital Twin-Scale Biomedical Models

Digital twin (DT) technologies are rapidly becoming a cornerstone of modern computational science and engineering workflows. From aerospace to manufacturing to healthcare, DTs offer a powerful way to simulate, monitor, and optimize complex systems by creating virtual replicas that evolve in real time with data. A pressing challenge in DT technologies is accurate and actionable information about uncertainty. We discuss advances in UQ methodology in practice in the context of cardiac digital twins (CDT), i.e., personalized models built from patient-specific clinical data, and virtual heart populations (VHP), considering heart function and outcomes across populations. These technologies require UQ tools that address model calibration, sensitivity analysis, and assessments of uncertainty that inform decision-making. Our procedures tackle high dimensionality of parameters involved in these models, epistemic uncertainty stemming from model errors, and finite population sampling errors, and we present translational software tools that we are building to address modeling and simulation of CDTs and VHPs, but that are also

applicable across broader DT domains.

Akil Narayan, Karli Gillette, Asher Merrill, Eli Child
University of Utah
akil@sci.utah.edu, karli.gillette@utah.edu, asher.merrill@utah.edu, eli.child@utah.edu

MS62

Promoting Confidence in Patient-Specific Finite Element Models of the Vascular System Through Uncertainty Quantification

Multiphysics-based simulations of the cardiovascular system are increasingly being integrated into clinical decision-making, surgical planning, medical device design, and regulatory affairs. We have incorporated advancements in uncertainty quantification (UQ) in a series of studies to evaluate the variability in the output of computational simulations of the fluid and solid coronary artery biomechanical environments due to model input uncertainty. Probability density functions were fit to population-derived or patient-specific values and sampled. Idealized and patient-specific coronary artery computational models of the fluid and solid domains were constructed and solved in the finite element software suite, FEBio, to predict fluid wall shear stress and tissue first principal stress. Parameter-to-model-output emulators were constructed via polynomial chaos expansion using the UQ software, UncertainSCI. Results highlight the variability in predicted mechanical metrics due to input parameter uncertainties, and particularly the impact of patient-specific versus population-derived data on model uncertainty. These findings underscore the role of UQ in evaluating model robustness and present uncertainty-aware frameworks to guide model refinement and improve reliability, interpretability, and clinical translation.

Lucas Timmins
Department of Biomedical Engineering
Texas A&M University
lucas.timmins@tamu.edu

MS63

Dealing with Uncertainties in Modeling Nuclear Reactions

We discuss the development of an uncertainty-quantified phenomenological model for nuclear reactions, a parameterization of the quantum-mechanical interaction between the nucleons and the nucleus, calibrated to many experimental data sets of angular cross sections. We primarily focus on the development of the likelihood model and the Bayesian calibration. We discuss the challenge of mitigating overconfident posteriors in the presence of experimental constraints consisting of many data points without provided covariances. We discuss treatment of correlated model discrepancy using a low-rank covariance structure in Legendre polynomial space over angles. We also discuss the novel computational tools developed, which allow for this massively-parallel calibration to be computationally tractable.

Kyle Beyer
Facility for Rare Isotope Beams
Michigan State University
beyerk@frib.msu.edu

MS63

Emulation of Stochastic Simulation Model Describing Heavy-Ion Collisions

The simulation of relativistic heavy-ion collisions is critical in understanding the evolution of Quark-Gluon Plasma during the early phases. In particular, the viscous anisotropic hydrodynamics (VAH) simulations provide an important component for transitioning smoothly across time. The Bayesian calibration of VAH is challenging due to multiple factors, posing a need for an efficient and accurate emulator. First, the VAH outputs are stochastic, in which the statistical variability can be studied via replications. Second, the VAH outputs are high-dimensional ($d \approx 100$), corresponding to potentially correlated observables. Third, one VAH simulation requires $\sim$ 1000 core-hours to complete, making direct calibration infeasible. We propose the use of a novel Gaussian process-based methodology to construct an effective emulator. The model makes use of a flexible representation of the underlying functional components, while handling the correlated outputs by selecting a data-driven transformation. The construction of emulator that accurately represents the uncertainty facilitates the calibration of the VAH simulations. Computational experiments show that the improved quality of the new method, measured by predictive accuracy, statistical coverage, and proper scoring rules.

Moses Chan, Evan Barnett
Northwestern University
moses.chan@northwestern.edu, evanbarnett2027@northwestern.edu

MS63

Model Combination and Orthogonalization for Controlled Extrapolations

Some of the most interesting applications of data from physics lies at the extremes. In nuclear physics, questions surrounding both the formation of everyday elements in the hearts of stars and in the most exotic and ephemeral systems produced only in terrestrial labs all rely on data that we sometimes have no way to access experimentally. This is where physicists turn to models to propagate the information in the experimentally accessible region off to the extremes. Of course, these models have flaws and shortcomings, though a substantial amount of effort has been spent in recent years to ground these predictions in principled uncertainty quantification frameworks. This talk will discuss one such approach to this issue by considering the whole zoo of models and predictions that have been made by experts over the years and combining them in a way that minimizes overlap between related models. This leads to an efficient UQ framework for multi-model combination that can be easily scaled to include more models and more predictions without a substantial cost in either the training or the inference stage.

Kyle Godbey
Facility for Rare Isotope Beams
Michigan State University
godbey@frib.msu.edu

MS63

Implicit Neural Representations for Solving Parametrized PDEs with Uncertainty Quantification: Towards Applications in Nuclear Theory

Implicit Neural Representations (INRs) offer a powerful framework for learning compact, continuous representations of a variety of signal types ranging from images to

solutions of partial difference equations (PDEs). This presentation explores how INRs can be leveraged to model high-dimensional solution spaces of parametrized PDEs without reliance on high fidelity solutions or dense coordinate grids for training data. By opting for flexible operator learning through hypernetworks, our INR approach can efficiently capture solutions across diverse parameter regimes, enabling scalable approximation of those PDE families. To address reliability in scientific applications, we incorporate Bayesian uncertainty quantification, focusing on Laplace approximation style methods to derive posterior distributions over network weights and propagate them into credible solution intervals. Additionally, we provide methods that adapt classical Laplace approximation techniques to training regimes where the loss landscape is not smooth or twice differentiable, accounting for nuanced physics-based constraints. As case studies, we examine the quantum harmonic oscillator and a nonlinear, non-affine system, showing how INRs capture eigenstates, integrate physics-informed losses, and handle challenging parameter dependencies. We conclude with directions for scaling to more complex systems, refining UQ methods, and extending applications in nuclear physics.

Sophia Zorek
Rice University
saz2@rice.edu

Kyle Godbey, Pablo Giuliani
Facility for Rare Isotope Beams
Michigan State University
godbey@frib.msu.edu, giulia27@msu.edu

Edgard Bonilla
Stanford University
edgard@stanford.edu

Tyler Bagwell, Frederi Viens
Rice University
teb6@rice.edu, fv15@rice.edu

MS64

Nonlinear Estimation of Cloud Parameters with Ensemble Conjugate Transform Inversion Methods

One of the major challenges in cloud and precipitation modeling is the presence of bounded quantities. For example, precipitation amounts are inherently non-negative and often exhibit highly skewed distributions near their lower bounds. Similar constraints apply to cloud microphysical parameters, which govern the evolution of cloud processes in Earth system models. As a result, parameter estimation methods based on linear-Gaussian assumptions frequently struggle to accurately calibrate these microphysical parameters from cloud-sensitive observations. In this talk, we investigate the feasibility of a newly developed inversion algorithm the Ensemble Conjugate Transform Inversion (ECTI) for estimating microphysical parameters within a column cloud model. ECTI builds on a recent generalization of Kalman filter theory to non-Gaussian distributions (Chipilski 2025: https://doi.org/10.1175/JAS-D-24-0171.1) and leverages invertible transport maps. We will highlight the advantages of ECTI implementations that incorporate normalizing flows, a class of generative AI methods, and compare their performance against the benchmark Ensemble Kalman Inversion (EKI) and other nonlinear generalizations.

Hristo G. Chipilski
Florida State University
hchipilski@fsu.edu

MS64

Latent Ensemble Score Filters for High-dimensional Data Assimilation with Sparse Observations

Accurate modeling and prediction of complex physical systems often rely on data assimilation techniques to correct errors inherent in model simulations. Traditional methods like the Ensemble Kalman Filter (EnKF) and its variants as well as the recently developed Ensemble Score Filters (EnSF) face significant challenges when dealing with high-dimensional and nonlinear Bayesian filtering problems with sparse observations, which are ubiquitous in real-world applications. In this paper, we propose a novel data assimilation method, Latent-EnSF, which leverages EnSF with efficient and consistent latent representations of the full states and sparse observations to address the joint challenges of high dimensionlity in states and high sparsity in observations for nonlinear Bayesian filtering. We introduce a coupled Variational Autoencoder (VAE) with two encoders to encode the full states and sparse observations in a consistent way guaranteed by a latent distribution matching and regularization as well as a consistent state reconstruction. With comparison to several methods, we demonstrate the higher accuracy, faster convergence, and higher efficiency of Latent-EnSF for two challenging applications with complex models in shallow water wave propagation and medium-range weather forecasting, for highly sparse observations in both space and time.

Phillip Si, Peng Chen
Georgia Institute of Technology
psi6@gatech.edu, pchen402@gatech.edu

MS64

Direct Assimilation of Observables

We propose a method for conducting data assimilation for data-driven models of unknown partially observed dynamical and PDE systems. When only a subset of the systems state variables is observed, memory-dependent models can account for the unobserved dynamics and provide surrogates for the evolution of the observable states. We perform assimilation in this reduced setting by leveraging memory-dependent deep neural networks based in flow-map learning and the MoriZwanzig formalism. Given noisy and partial observations of the observable states over an assimilation window, we aim to approximate the corresponding memory trajectory to update the model's prediction. This assimilation enables the model to generate accurate long-time forecasts of the observable states, including for unseen initial conditions. In addition, we investigate model-correction of the surrogate during assimilation to further improve our approximation of the unknown governing dynamics.

Caroline Tatsouka, Dongbin Xiu
Ohio State University
tatsouka.3@osu.edu, xiu.16@osu.edu

Youssef M. Marzouk
Massachusetts Institute of Technology

ymarz@mit.edu

MS64

Diffusion-Based Supervised Learning of Generative Models for Efficient Sampling of Multimodal Distributions

We propose a hybrid generative model for efficient sampling of high-dimensional, multimodal probability distributions for Bayesian inference. Traditional Monte Carlo methods, such as the Metropolis-Hastings and Langevin Monte Carlo sampling methods, are effective for sampling from single-mode distributions in high-dimensional spaces. However, these methods struggle to produce samples with the correct proportions for each mode in multimodal distributions, especially for distributions with well separated modes. To address the challenges posed by multimodality, we adopt a divide-and-conquer strategy. We start by minimizing the energy function with initial guesses uniformly distributed within the prior domain to identify all the modes of the energy function. Then, we train a classifier to segment the domain corresponding to each mode. After the domain decomposition, we train a diffusion-model-assisted generative model for each identified mode within its support. Once each mode is characterized, we employ bridge sampling to estimate the normalizing constant, allowing us to directly adjust the ratios between the modes. Our numerical examples demonstrate that the proposed framework can effectively handle multimodal distributions with varying mode shapes in up to 100 dimensions. An application to Bayesian inverse problem for partial differential equations is also provided.

Hoang A. Tran
Oak Ridge National Laboratory
Computer Science and Mathematics Division
tranha@ornl.gov

Feng Bao
Florida State University
Department of Mathematics
fbao@fsu.edu

Guannan Zhang, Dan Lu
Oak Ridge National Laboratory
zhangg@ornl.gov, lud1@ornl.gov

Zezhong Zhang
Auburn University
zez0002@auburn.edu

MS65

Scaling Up UQ: From Algorithm Development to Supercomputing

UQ for compute-heavy simulations requires a combination of advanced UQ algorithms, efficient surrogates and scalable simulators. This talk presents (1) recent advances in algorithms (e.g. parallel prefetching for multilevel UQ), (2) software frameworks allowing even UQ prototypes to transparently control supercomputers (UM-Bridge, Tesseract) and (3) a large-scale UQ application that combines the best of the three worlds to learn earthquake model parameters from full-field seismic simulations and indirect observations.

Linus Seelinger
Karlsruhe Institute of Technology (KIT)
Scientific Computing Center
mail@linusseelinger.de

MS65

Learning Low-Rank Neutron Transport Dynamics on Linear and Nonlinear

In this presentation, we propose semi-intrusive low-rank model reduction strategies for time-dependent neutron transport problems. Starting from a full-order discretization of the transport equation in space and angle, we reformulate the dynamics in global monolithic form and construct reduced models using both linear subspaces and nonlinear manifolds approaches. The nonlinear framework leverages quadratic lifting maps (element-wise and tensorial) that embed low-dimensional latent variables into higher-dimensional reduced spaces, enhancing the expressive power of the reduced basis. To avoid the prohibitive cost of projecting transport operators, we introduce a data-driven operator inference strategy that directly estimates reduced operators from high-fidelity simulation data. Numerical experiments on a one-dimensional slab benchmark show that tensorial lifting achieves markedly improved accuracy compared to linear and element-wise quadratic approaches while maintaining significant computational speed-ups. The results highlight the potential of nonlinear manifold reduction combined with operator inference to deliver efficient and accurate surrogate models for neutron transport dynamics, with promising applicability to large-scale, multi-dimensional settings.

Francesco A. B. Silva
Texas A&M Engineering Experiment Station
Nuclear Engineering
francesco.silva@tamu.edu

MS66

Modeling of Spatial Asymmetries in Temperatures in North America with Normalizing Flows

Combining strengths from deep learning and classical statistics can help describe complex relationships between variables that can have significant impacts (e.g., environmental, or financial applications). Neural networks learn complicated nonlinear relationships from large datasets under limited parametric assumptions. By definition, the number of samples of rare events is small, which limits the ability of the data-hungry, nonparametric neural network to describe these events. Inspired by recent cold winter weather events in North America caused by atmospheric blocking, we examine several probabilistic generative models for the entire multivariate probability distribution of daily boreal winter surface air temperature. Compared to vine copulas, the statistical standard for multivariate copula modeling, deep learning methods show improved ability to reproduce complicated asymmetries in the spatial distribution of ERA5 temperature reanalysis, including the spatial extent of in-sample rare events.

Julie Bessac
US National Renewable Energy Laboratory
julie.bessac@nrel.gov

Mitchell Krock
University of Missouri

mk52n@missouri.edu

MS66

Generative Neural Networks towards Intractable Bayesian Inverse Problems

In this talk we consider uncertainty quantification (UQ) for inverse problems applied to physical systems governed by differential equations that additionally have stochasticity. Each "dimension of complexity" - the forward problem, inverse problem, UQ - dramatically increases the computational challenges. Even simple forward problems become challenging when framed as inverse problems if the data misfit term generates highly nonlinear and/or nondifferentiable mappings. Extending inference to UQ adds another dimension of challenges and leads to computational intractability for complex systems. This talk investigates whether generative neural networks can break some of the complexity. The goal is for a neural network (NN) to represent an "inverse mapping": from observational data to parameters. State-of-the-art NNs can be trained to predict point estimates or empirical representations of distributions, and as such delivering UQ. Using NNs in the context of inverse mappings is also referred to as amortized inference; and the literature also classifies inverse problems and UQ (with or without NNs) as simulation-based inference. The talk considers open challenges remaining at the intersection of inverse problems, deep learning, and computational mathematics: Which NN expressivity can approximate solutions of ill-posed inverse problems? What is the impact of finite training data and noise? How do generative NNs scale as sampling generators for Bayesian inverse problems?

Johann Rudi
Virginia Tech
jrudi@vt.edu

Deep Ray
University of Maryland, College Park
deepray@umd.edu

Anwesa Dey
University of Utah
deyanwesa@gmail.com

Andreas Mang
University of Houston
Department of Mathematics
andreas@math.uh.edu

Ionut-Gabriel Farcas
Virginia Tech
farcasi@vt.edu

MS66

Generative Learning for Forward and Inverse Modeling of Complex Spatiotemporal Physics

Generative models offer a principled way to represent uncertainty in high-dimensional dynamical systems, yet most surrogates in scientific computing remain point predictors tied to a single task or regime. This talk presents a hybrid physics-integrated generative framework that unifies forward forecasting, flow reconstruction/super-resolution, data assimilation, and inverse modeling for complex spatiotemporal PDEs. The core idea is to learn a probabilistic prior over fields/trajectories using diffusion or conditional flow-matching, and to couple that prior with a coarse-grained differentiable solver during inference via solver-conditioned diffusion posterior sampling. Practically, we model either (i) small-scale residuals/SGS stresses conditioned on the coarse state, or (ii) the full field, and enforce incompressibility, energy/enstrophy budgets, spectral targets, and symmetry through architectural and loss constraints. These yield calibrated ensembles that respect physics while accommodating chaotic phase drift and partial/noisy observations. Across homogenous, isotropic to wall-bounded turbulence testbeds, this yields high-k spectral recovery and reduced long-horizon budget drift, delivers well-calibrated posteriors for state and parameter estimation, and enables zero-shot data assimilation and inverse modeling.

Meet Parikh, Jian-Xun Wang
Cornell University
mhp66@cornell.edu, jw2837@cornell.edu

MS66

Hierarchical Flow Matching with Mini-Batch Couplings

Flow matching has emerged as a compelling generative modeling approach that is widely used across domains. During training, flow matching learns to model a velocity field. At inference, to generate samples, an ordinary differential equation (ODE) is numerically solved via forward integration of the modeled velocity field. To better capture the multi-modality that is inherent in typical velocity fields, we introduce hierarchical flow matching. It uses a hierarchy of ODEs that are numerically integrated when generating data. Each level of the hierarchy of ODEs captures the distribution of the next level, just like vanilla flow matching uses the velocity field to capture a multi-modal data distribution. While this hierarchy enables modeling multi-modal distributions at any hierarchy level, the complexity of the modeled distributions remains identical across levels of the hierarchy. We therefore further investigate how to progressively adjust the complexity of the distributions across different levels of the hierarchy via mini-batch couplings. We show the benefits of mini-batch couplings in hierarchical flow matching via compelling results on synthetic and imaging data, and provide theoretical justification.

Zhizhen Zhao
University of Illinois
zhizhenz@illinois.edu

MS67

Optimal Control for Transformer Architectures: Enhancing Generalization, Robustness and Efficiency

We study Transformers through the perspective of optimal control theory, using tools from continuous-time formulations to derive actionable insights into training and architecture design. This framework improves the performance of existing Transformer models while providing desirable theoretical guarantees, including generalization and robustness. Our framework is designed to be plug-and-play, enabling seamless integration with established Transformer models and requiring only slight changes to the implementation. We conduct seven extensive experiments on tasks motivated by text generation, sentiment analysis, image classification, and point cloud classification. Experimental results show that the framework improves the test performance of the baselines, while being

more parameter-efficient. On character-level text generation with nanoGPT, our framework achieves a 46% reduction in final test loss while using 42% fewer parameters. On GPT-2, our framework achieves a 5.6% reduction in final test loss, demonstrating scalability to larger models. To the best of our knowledge, this is the first work that applies optimal control theory to both the training and architecture of Transformers. It offers a new foundation for systematic, theory-driven improvements and moves beyond costly trial-and-error approaches.

Kelvin Kan
Department of Mathematics, Emory University
kaifung0620@gmail.com

MS67

Transformer As a Nonlinear Predictor

I will describe a mathematical framework for causal nonlinear prediction in settings where observations are generated from an underlying hidden Markov model (HMM). Both the problem formulation and the proposed solution are motivated by the decoder-only transformer architecture, in which a finite sequence of observations (tokens) is mapped to the conditional probability of the next token. I will describe these correspondences as part of my talk.

Prashant G. Mehta
Department of Mechanical Science and Engineering
University of Illinois at Urbana Champaign
mehtapg@illinois.edu

MS67

Activation Steering with a Feedback Controller

Controlling the behaviors of large language models (LLM) is fundamental to their safety alignment and reliable deployment. Yet, existing steering methods are primarily driven by empirical insights and lack theoretical performance guarantees. In this work, we develop a control-theoretic foundation for activation steering by showing that popular steering methods correspond to the proportional (P) controllers, with the steering vector serving as the feedback signal. Building on this finding, we propose ProportionalIntegralDerivative (PID) Steering, a principled framework that leverages the full PID controller for activation steering in LLMs. The proportional (P) term aligns activations with target semantic directions, the integral (I) term accumulates errors to enforce persistent corrections across layers, and the derivative (D) term mitigates overshoot by counteracting rapid activation changes. This closed-loop design yields interpretable error dynamics and connects activation steering to classical stability guarantees in control theory. Moreover, PID Steering is lightweight, modular, and readily integrates with state-of-the-art steering methods. Extensive experiments across multiple LLM families and benchmarks demonstrate that PID Steering consistently outperforms existing approaches, achieving more robust and reliable behavioral control.

Tan Nguyen, Dung Nguyen
National University of Singapore
tanmn@nus.edu.sg, dungnv@u.nus.edu

Nhi Pham
Vin University
nhiiyennphamm@gmail.com

Hieu Vu, Lei Zhang
National University of Singapore
vmhieu17@gmail.com, matzhlei@nus.edu.sg

MS67

Understanding in-Context Learning on Structured Manifolds: Bridging Attention to Kernel Methods

While in-context learning (ICL) has achieved remarkable success in natural language and vision domains, its theoretical understandingparticularly in the context of structured geometric dataremains unexplored. This paper initiates a theoretical study of ICL for regression of Holder functions on manifolds. We establish a novel connection between the attention mechanism and classical kernel methods, demonstrating that transformers effectively perform kernel-based prediction at a new query through its interaction with the prompt. This connection is validated by numerical experiments, revealing that the learned queryprompt scores for Holder functions are highly correlated with the Gaussian kernel. Building on this insight, we derive generalization error bounds in terms of the prompt length and the number of training tasks. When a sufficient number of training tasks are observed, transformers give rise to the minimax regression rate of Holder functions on manifolds, which scales exponentially with the intrinsic dimension of the manifold, rather than the ambient space dimension. Our result also characterizes how the generalization error scales with the number of training tasks, shedding light on the complexity of transformers as in-context kernel algorithm learners. Our findings provide foundational insights into the role of geometry in ICL and novels tools to study ICL of nonlinear models.

Zhaiming Shen
Georgia Institute of Technology
zshen49@gatech.edu

MS68

Invariant Measures for Data-Driven Dynamical System Identification

In this talk, we propose a novel approach for performing dynamical system identification, based upon the comparison of simulated and observed physical invariant measures. While standard methods directly treat time-trajectories as inference data, we instead seek models fitting the observed global time-invariant statistics. With this change in perspective, we gain robustness against pervasive challenges in system identification including noise, chaos, and slow sampling. In the first half of the talk, we pose the system identification task as a partial differential equation (PDE) constrained optimization problem, in which synthetic stationary solutions of the Fokker-Planck equation, obtained as fixed points of a finite-volume discretization, are compared to physical invariant measures extracted from observed trajectory data. In the latter half, we use Takens' embedding theory to introduce a critical data-dependent coordinate transformation which can guarantee unique system identifiability from the invariant measure alone. This contribution resolves a major challenge of system identification through invariant measures, as systems exhibiting distinct transient behaviors may still share the same time-invariant statistics in their state-coordinates. Throughout, we present comprehensive numerical tests which highlight the effectiveness of our approach on a variety of challenging system identification tasks.

Jonah Botvinick-Greenhouse, Yunan Yang
Cornell University

jrb482@cornell.edu, yunan.yang@cornell.edu

MS68

Transformers Through the Lens of Support-preserving Maps Between Measures

Transformers are deep architectures that define "in-context maps" which enable predicting new tokens based on a given set of tokens. In this work, we model neural networks as maps on probability measures and study the question "what kind of maps between measures are transformers?" We fully characterize the properties of maps between measures that enable them to be represented in terms of in-context maps via a push forward. On the one hand, these include transformers; on the other hand, transformers universally approximate representations with any continuous in-context map. These properties are preserving the cardinality of support and that the regular part of their Fréchet derivative is uniformly continuous. Moreover, we show that the solution map of the Vlasov equation, which is of nonlocal transport type, for interacting particle systems in the mean-field regime for the Cauchy problem satisfies the conditions on the one hand and, hence, can be approximated by a transformer; on the other hand, we prove that the measure-theoretic self-attention has the properties that ensure that the infinite depth, mean-field measure-theoretic transformer can be identified with a Vlasov flow.

Takashi Furuya
Doshisha University
tfuruya@mail.doshisha.ac.jp

MS68

The Measure Evolution of Hyperbolic PDEs with Random Data

We study pointwise statistics in 1D hyperbolic PDEs with random initial data–that is, the distribution of the solution at any given point in space and time. For linear problems, the PDF evolves via the same PDE as the original system. In the nonlinear case, we derive a linear transport equation for the CDF, valid up to shock formation, beyond which a counterexample shows no such equations exist. For systems, we extend this to multi-point statistics. These equations offer a computationally efficient alternative to Monte Carlo methods and surrogate models. They also yield a priori statistical error bounds for standard Monte Carlo methods.

Amir Sagiv
NJIT
amir.sagiv@njit.edu

MS69

Enhancing Supervised Dimension Reduction via Multi-Model Data Fusion

Quantities of interest (QoIs) in large scientific and engineering applications are typically not independent, and can therefore be represented on a manifold, i.e. with fewer dimensions than in the original ambient space. In this context, strategies for linear dimensionality reduction (DR), like Principal Component Analysis (PCA), are commonly employed to identify linear transformations of the input parameters into a more parsimonious coordinate system, still accurately capturing most of the variability in the quantities of interest (QoIs). Alternatively, strategies for nonlinear DR (like diffusion maps and autoencoders) represent high-dimensional data on nonlinear low-dimensional manifolds. In a recent contribution [1], we introduced NeurAM, a data-driven paradigm where a nonlinear reduced-order manifold is discovered that preserves the input-to-output map. In this talk, we extend this idea by formulating a supervised loss in terms of a multi-fidelity surrogate model, as opposed to a single-fidelity surrogate as in NeurAM. On multiple test problems with varying complexity, we demonstrate that this extension leads to estimators of reduced variance, while also reducing the amount of required high-fidelity training data. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525 [1] Zanoni, et al., NeurAM: nonlinear dimensionality reduction for uncertainty quantification through neural active manifolds, arXiv:2408.03534, 2024.

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Owen N. Davis
University of New Mexico
ondavis@sandia.gov

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

Andrea Zanoni
Centro di Ricerca Matematica
Scuola Normale Superiore, Pisa, Italy
andrea.zanoni@sns.it

MS69

Multilevel Monte Carlo for Bayesian Optimization and Deep Learning

Bayesian methods provide a principled foundation for building reliable, uncertainty-aware AI systems. We apply the powerful multilevel Monte Carlo (MLMC) approach to problems formulated in a Bayesian framework. First, we apply MLMC to the problem of Bayesian Optimization (MLMCBO), in order to reduce the complexity of lookahead acquisition functions and yield faster global optimization of expensive black-box functions. Second, we apply MLMC to Bayesian deep learning with convergent trace-class neural networks (TNN) and MLSMC samplers, thereby enabling more efficient quantification of uncertainty.

Kody Law
University of Manchester
kodylaw@gmail.com

MS69

Multi-fidelity Ensembles for UQ in Monte Carlo Transport Applications

Uncertainty propagation is a key component of predictive science. Algorithms that scale efficiently for large dimensions have been proposed, however, dealing high-fidelity applications still remains challenging due to their computational cost, which limits the available number of realizations. For instance, Monte Carlo radiation transport solvers are notoriously expensive to converge to satisfactory statistical levels with limited number of particles. Multi-fidelity strategies were introduced to lower the computational burden of uncertainty analysis: the idea is to combine limited high-fidelity realizations with a larger set of

realizations obtained from lower fidelity models. Sampling methods exploit correlation among models to reduce the estimator variance. Unfortunately, limited studies exist on the effect that the stochasticity intrinsic in solvers or physical processes, e.g., random media, exercise on the correlation/cost trade-off among models. Building on our recent modeling efforts in representing stochastic media for radiation transport applications, we explore the cost/correlation trade-off that models, their numerical approximations, and their implementations can provide to multi-fidelity studies in the context of radiation transport applications when stochastic media effects are important. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525. This work is supported by SNL LDRD Project 237058.

Alec Shelley
Stanford University
ams01@stanford.edu

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Aaron Olson
Sandia National Laboratories
aolson@sandia.gov

Daniel M. Tartakovsky
Stanford University
tartakovsky@stanford.edu

MS70

Multilevel Lattice-based Kernel Approximation for Random Elliptic PDEs

Kernel interpolation using lattice rules have shown to be efficient methods for approximating solutions to elliptic PDE with periodic random coefficients. The two key features of kernel methods with lattice rules are that the periodicity gives rise to higher-order convergence and the lattice structure can be exploited for fast computation using the FFT. Building upon the work of Kaarnioja, Kazashi, Kuo, Nobile, Sloan (Numer. Math. 2022), this talk outlines how multilevel methods can be leveraged to further improve the efficiency. I will discuss the complexity analysis of the multilevel kernel approximation, which relies on new mixed regularity bounds for the periodic setting. Finally, I will present numerical results demonstrating the improved efficiency compared so single level kernel interpolation.

Alexander D. Gilbert
School of Mathematics and Statistics
University of New South Wales
alexander.gilbert@unsw.edu.au

MS70

Fast Multitask Gaussian Process Regression

Gaussian process (GP) regression is a powerful probabilistic modeling technique with built-in uncertainty quantification. When one has access to multiple correlated simulations (tasks), it is common to fit a multitask GP (MTGP) surrogate which is capable of capturing both inter-task and intra-task correlations. However, with a total of N evaluations across all tasks, fitting an MTGP is often infeasible due to the $\mathcal{O}(N^2)$ storage and $\mathcal{O}(N^3)$ computations required to store, solve a linear system in, and compute the determinant of the $N \times N$ Gram matrix of pairwise kernel evaluations. In the single-task setting, one may reduce the required storage to $\mathcal{O}(N)$ and computations to $\mathcal{O}(N \log N)$ by fitting "fast GPs" which pair low-discrepancy design points from quasi-Monte Carlo to special kernel forms which yields nicely structured Gram matrices, e.g., circulant matrices. This article generalizes fast GPs to fast MTGPs which pair low-discrepancy design points for each task to special product kernel forms which yields nicely structured block Gram matrices, e.g., circulant block matrices. An algorithm is presented to efficiently store, invert, and compute the determinant of such Gram matrices with optionally different sampling nodes and different sample sizes for each task. A GPU-compatible, open-source Python implementation is made available in the FastGPs package (https://alegresor.github.io/fastgps/). We validate the efficiency of our algorithm and implementation compared to standard techniques on a range of problems with low numbers of tasks and large sample sizes.

Aleksei Sorokin
Illinois Institute of Technology
asorokin@hawk.iit.edu

Pieterjan Robbe
Sandia National Laboratories
pmrobbe@sandia.gov

Fred J. Hickernell
Illinois Institute of Technology
Dept of Applied Math & Ctr for Interdisc Sci Comput
hickernell@illinoistech.edu

MS72

Capturing Closure Terms in Nonlinear Dynamics Using Data Assimilation

Constitutive models play a crucial role in the analysis of dynamical systems in numerous engineering and scientific disciplines. These models provide a systematic and quantitative framework for (a) understanding the behavior, and (b) predicting the evolution of complex systems. Oftentimes such models, by design, overlook certain aspects of the underlying physics for a variety of reasons, such as the need to reduce computational requirements to simulate those systems or the lack of complete understanding of complex physical phenomenon. As models continue to be used in exploring new regimes of operation, previously overlooked and unresolved dynamics start to gain significance due to declining predictive accuracies. In order to maintain a target level of predictive accuracy in such settings, it is important to capture closure terms using data-driven approaches. We promote the use of stochastic dynamical models to capture missing closure terms, to be trained within a data assimilation setting. This talk will examine the utility in applying advanced Bayesian model selection techniques to answer questions related to the type of stochastic model needed to capture closure terms. For numerical investigation, we consider a slight modification of the nonlinear Duffing oscillator to examine the efficacy of the proposed framework. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Mohammad Khalil, Eric Crislip, Rileigh Bandy, Teresa Portone, Kyle Neal
Sandia National Laboratories
mkhalil@sandia.gov, escrisl@sandia.gov,
rjbandy@sandia.gov, tporton@sandia.gov,

kneal@sandia.gov

MS72

On Over-Parametrized Models and Sobolev Training

With Sobolev training, neural networks are provided data about both the function of interest and its derivatives. This setting is prevalent in scientific machine learningappearing in molecular dynamics emulators, derivative-informed neural operators, and predictors of summary statistics of chaotic dynamical systemsas well as in traditional machine learning tasks like teacher-student model distillation. However, fundamental questions remain: How does overparameterization influence performance? What role does the signal-to-noise ratio play? And is additional derivative data always beneficial? In this work, we study these questions using tools from statistical physics and random matrix theory. In particular, we consider Sobolev training in the proportional asymptotics regime in which the problem dimensionality d, single hidden-layer features p, and training points n grow to infinity at fixed ratios. We focus on target functions modeled as single-index models (i.e., ridge functions with a single intrinsic dimension), providing theoretical insights into the effects of derivative information in high-dimensional learning.

Matt Li
Massachusetts Institute of Technology
mtcli@mit.edu

Katharine Fisher
MIT
kefisher@mit.edu

Timo Schorlepp
Courant Institute of Mathematical Sciences
New York University
timo.schorlepp@nyu.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS72

Operator Learning for Surrogate Modeling in Data Assimilation

Scientific machine learning is transforming the field of data assimilation by providing accurate and cheap-to-evaluate surrogate models for many dynamical systems. In the first part of this talk, I will show that Fourier Neural Operators (FNOs) enjoy polynomial sample complexity for estimating the solution map of a wide class of evolution PDEs. Our theory leverages that these PDEs can be accurately approximated via pseudo-spectral methods, and that these numerical schemes can in turn be efficiently learned using FNOs. In the second part of this talk, I will show that ensemble Kalman filters with machine-learned surrogate models achieve long-time filter accuracy under standard observability assumptions on the true dynamics and observation models.

Nisha Chandramoorthy
The University of Chicago
nishac@uchicago.edu

Daniel Sanz-Alonso, Nathan Waniorek
University of Chicago
sanzalonso@uchicago.edu, waniorek@uchicago.edu

MS73

Extreme Event Aware Learning

Quantifying and predicting rare and extreme events persists as a crucial yet challenging task in understanding complex dynamical systems. Many practical challenges arise from the infrequency and severity of these events, including the considerable variance of simple sampling methods and the substantial computational cost of high-fidelity numerical simulations. Numerous data-driven methods have recently been developed to tackle these challenges. However, a typical assumption for the success of these methods is the occurrence of multiple extreme events, either within the training dataset or during the sampling process. To overcome this limitation, we introduce the framework of Extreme Event Aware or η-learning which does not assume the existence of extreme events in the available data. η-learning reduces the uncertainty even in 'unchartered' extreme event regions, by enforcing the extreme event statistics of a few observables during training, which can be available or assumed through qualitative arguments or estimated with unlabeled data. This type of statistical regularization results in models that fit the observed data and align with the prescribed observable statistics, enabling the generation of unprecedented extreme events even when the training data lack extremes therein. Theoretical results based on optimal transport as well as extensive numerical experiments are offered to justify and illustrate the favorable properties of the η-learning framework.

Kai Chang
Massachusetts Institute of Technology
kaichang@mit.edu

Themistoklis Sapsis
Massachusetts Institute of Techonology
sapsis@mit.edu

MS73

Generative Prediction-Correction Framework for Long-Time Emulations of Spatially-Resolved Climate Extremes

Accurately quantifying the increased risks of climate extremes requires generating large ensembles of climate realization across a wide range of emissions scenarios, which is computationally challenging for conventional Earth System Models. We propose GEN2, a generative prediction-correction framework for an efficient and accurate forecast of the extreme event statistics. The prediction step is constructed as a conditional Gaussian emulator, followed by a non-Gaussian machine-learning (ML) correction step. The ML model is trained on pairs of the reference data and the emulated fields nudged towards the reference, to ensure the training is robust to chaos. We first validate the accuracy of our model on historical ERA5 data and then demonstrate the extrapolation capabilities on various future climate change scenarios. When trained on a single realization of one warming scenario, our model accurately predicts the statistics of extreme events in different scenarios, successfully extrapolating beyond the distribution of training data.

Themistoklis Sapsis
Massachusetts Institute of Techonology
sapsis@mit.edu

Mengze Wang, Benedikt Barthel Sorensen
Massachusetts Institute of Technology
mzwang@mit.edu, bbarthel@mit.edu

MS73

Scalable Extreme Event Estimation for Stochastic Differential Equations Via Large Deviation Theory and Carleman-Fredholm Determinants

Estimating the probability of extreme events is an important problem in many scientific and engineering disciplines, as they are often associated with rare system failures or catastrophes. Asymptotically, rare event probabilities can be estimated with a Laplace approximation, known as (precise) large deviation theory. The approach involves (i) solving an optimization problem to find the most likely realization of the random parameter leading to a prescribed outcome, and (ii) calculating a determinant to account for Gaussian perturbations around the minimizer. In this talk, I will discuss how to carry out these steps numerically in a scalable way for extreme events in stochastic differential equations (SDEs) and stochastic partial differential equations (SPDEs) with Brownian noise. In particular, I will highlight the necessity to treat the determinant calculation in step (ii) correctly from an infinite-dimensional point of view to ensure scalability of the numerical method to high dimensions. This leads to either a Fredholm or Carleman-Fredholm determinant computation, depending on whether the second variation of the noise-to-event map is trace-class or only Hilbert-Schmidt. To illustrate these points, I will consider two examples of extreme event estimates for S(P)DEs with multiplicative noise: a predator-prey SDE as a toy problem, and a 2D random advection-diffusion problem as a nontrivial SPDE. Presentation based on: arXiv:2502.20114 (joint work with T. Grafke).

Timo Schorlepp
Courant Institute of Mathematical Sciences
New York University
timo.schorlepp@nyu.edu

Tobias Grafke
Warwick Mathematics Institute
University of Warwick
T.Grafke@warwick.ac.uk

MS74

Posterior Covariance Structures in Gaussian Processes

In this talk, we present a comprehensive analysis of the posterior covariance field in Gaussian processes, with applications to the posterior covariance matrix. Our geometric analysis reveals how the Gaussian kernels bandwidth parameter and the spatial distribution of the observations influence the posterior covariance as well as the corresponding covariance matrix, enabling straightforward identification of areas with high or low covariance in magnitude. Drawing inspiration from the a posteriori error estimation techniques in adaptive finite element methods, we also propose several estimators to efficiently measure the absolute posterior covariance field, which can be used for efficient covariance matrix approximation and preconditioning.

Difeng Cai
Southern Methodist University
ddcai@smu.edu

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

Yuanzhe Xi
Emory University
yuanzhe.xi@emory.edu

MS74

Uncertainty-Quantified X-Ray Ptychography Via Generative Prediction-Powered Inference

We propose a novel prediction-powered inference (PPI) framework for generative modeling in X-ray ptychographic imaging. Inverse problems in X-ray imaging are inherently ill-posed and plagued by noise, incomplete measurements, and model mismatch. Instead of relying solely on deterministic reconstructions, our method leverages a generative prior and formulates posterior recovery as an entropic optimal transport problem. By coupling generative sampling with PPI, we obtain calibrated UQ that captures multimodal solution landscapes and produces interpretable posterior diagnostics. This generative PPI approach provides explicit measures of reconstruction reliability, enabling more robust and transparent decision-making in X-ray imaging applications.

Huan He
Auburn University
huan.he@auburn.edu

MS74

Domain-Aware Gaussian Process for Quantitative Learning in Scientific Computing

Gaussian Process (GP) plays a significant role in the Bayesian statistical learning framework. It is a powerful tool for black-box optimization and for uncertainty quantification in digital twins. We have been developing a number of techniques to increase the usability of GP in large scale applications involving millions of data points. These techniques include fast kernel matrix inversion, new domain-aware kernels, multi-task, multi-objective and multi-fidelity learning. We will demonstrate the versatility of the GP framework through its use in parameter tuning of parallel programs, hardware knobs tuning of scientific instruments, and decision support for solving PDEs.

Xiaoye S. Li
Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS74

Efficient Optimization of Plasma Radiation Detector Configurations Using Imperfect Inference Models

The configurations of instruments fielded on an experiment affect the amount of information captured and the quality of subsequent inference. We investigate the problem of optimizing plasma x-ray radiation detectors in a magneto-inertial fusion experiment at Sandia National Laboratories. It is impossible to directly measure properties such as the temperature of the thermonuclear fusion plasma produced in these experiments because of the extreme environment and destructive nature of the experiment. Among other diagnostics, several detectors are placed with significant

standoff from the fusion target to capture the x-rays emitted by the fusion plasma, which can be used to infer some of its properties. To optimize the configuration of these detectors, a high-fidelity model (HFM) is used for simulating outputs and a low-fidelity model (LFM) is used for inference. We develop methods based on A- and L-optimality criteria that are efficient to compute while explicitly accounting for the discrepancy between the HFM and the LFM. The method allows us to find detector configurations that perform similarly to or better than the configuration obtained using an existing sampling-based optimization method while decreasing computational time by a factor of 50.

Difan Song
Georgia Institute of Technology
dfsong@gatech.edu

MS75

Structure in Dynamics-Encoded Deep Learning for Parameter Estimation

Incorporating a priori physics knowledge into scientific machine learning leads to more robust and interpretable algorithms. In this talk, we explore how structural characteristics of dynamics-embedded deep learning techniques affect resulting estimates of network and system parameters in dynamical systems applications.

Andrea Arnold
Worcester Polytechnic Institute
anarnold@wpi.edu

MS75

VPBoost: Accelerating Gradient Boosting with Variable Projection

Gradient boosting (GB) constructs predictive models by sequentially training weak learners that serve as descent directions in function space. While highly successful for decision trees, particularly through frameworks such as XGBoost, applying GB to neural networks (NNs) remains computationally expensive due to the high-dimensional, non-convex optimization required for each weak learner. Recent advances demonstrate that boosted NNs can outperform decision trees on complex data, yet training efficiency remains a critical bottleneck. Variable projection (VarPro) has emerged as a powerful optimization strategy for separable NNs, which decompose into nonlinear feature extraction followed by linear mapping. By analytically eliminating the linear weights through partial optimization, VarPro reduces the parameter space and accelerates convergence. We propose VPBoost, a novel algorithm that combines VarPro with an XGBoost's algorithm for gradient boosting. For each weak learner, VPBoost solves a reduced optimization problem over only the nonlinear parameters while maintaining optimal linear weights analytically. This exploitation of separability yields a bi-level optimization framework where the outer problem minimizes over featurizer parameters and the inner problem provides closed-form solutions for linear weights. Our methodology establishes convergence guarantees analogous to standard gradient descent while achieving faster training per weak learner and requiring fewer ensemble members. Moreover, gradient boosting integrates naturally with VarPro, reducing the computational and implementation complexity barrier that VarPro traditionally presents. Numerical experiments demonstrate substantial computational savings and improved generalization compared to traditional gradient boosting approaches for NNs, positioning VPBoost as an efficient alternative for scientific machine learning and data-intensive applications.

Abhijit Chowdhary
North Carolina State University
abhijit.chowdhary@tufts.edu

Elizabeth Newman
Emory University
e.newman@tufts.edu

Deepanshu Verma, Deepanshu Verma
Clemson University
dverma@clemson.edu, dverma@clemson.edu

MS75

Machine Learning for UQ in Inverse Problems Governed by Non-linear Dynamical Systems

We study neural networks (NNs) for estimating hidden parameters and quantifying uncertainty from noisy data. Parameter estimation is formulated as a Bayesian inverse problem for parametrized nonlinear systems. The problem poses major challenges for classical methods due to strong nonlinearities, nonconvexity, and sharp gradients. We show that NNs address these challenges by learning reconstruction maps from observational data. Using a single forward pass, the NN infers model parameters, noise parameters for both autocorrelated additive noise and stochastic differential equationdriven noise, and the posterior covariance matrix. We compare different NN architectures, analyze the impact of noise on accuracy, and report training times on specialized hardware. Our results demonstrate that NNs provide an efficient and flexible framework for parameter and uncertainty estimation in dynamical systems.

Pegah Amiri
University of Houston
pamiri@CougarNet.UH.EDU

German Villalobos
Department of Mathematics
University of Houston
gvillalobos@uh.edu

Johann Rudi
Virginia Tech
jrudi@vt.edu

Andreas Mang
University of Houston
Department of Mathematics
andreas@math.uh.edu

MS76

UncertaintyQuantification.jl: Reliability Analysis and Simulation under Uncertainty in Julia

Uncertainty quantification (UQ) is central to making reliable predictions in science and engineering, particularly when models are exposed to variability in inputs, parameters, or data. This talk presents UncertaintyQuantification.jl, a Julia package designed to streamline uncertainty modelling and reliability analysis for the computational sciences. The package implements some state-of-the-art reliability methods, including Monte Carlo, line sampling, and subset simulation, making it possible to efficiently esti-

mate rare event probabilities and failure domains in high-dimensional systems. Beyond more typical probabilistic approaches, it offers tools for handling imprecise probabilities, such as interval and set-valued uncertainty models, enabling a conservative quantification even when information is limited or ambiguous. The package integrates easily with external solvers and scales from desktop prototyping to HPC clusters. Numerical examples will illustrate how these methods may be used in challenging risk assessments involving stochastic variability and epistemic uncertainty, including applications to large-scale nuclear fusion models on HPC systems.

Ander Gray
UK Atomic Energy Authority (UKAEA)
ander.gray@ukaea.uk

Jasper Behrensdorf
Institute for Risk and Reliability
Leibniz Universität Hannover
behrensdorf@irz.uni-hannover.de

Lukas Fritsch, Matteo Broggi
Institute for Risk and Reliability
Leibniz University Hannover
fritsch@irz.uni-hannover.de, broggi@irz.uni-hannover.de

Michael Beer
Leibniz University Hannover
beer@irz.uni-hannover.de

MS76

Operator Learning for Wildfire and Smoke Dynamics

High-consequence physical systems often feature quite complex dynamics, which are too computationally burdensome within outer-loop analysis. Wildfires are a prominent example of such a physical, causing widespread health and financial damage each year. Typically, however, they are simulated with a two-way coupled system of partial differential equations (PDEs) for the fire front and the atmosphere. Surrogate modeling replaces a complex model with a less accurate but significantly faster approximation, which makes repeated querying tractable. The proper choice of surrogate depends not only on the properties of the complex model but also on the desired outer-loop analysis. In this talk, we discuss two outer-loop analyses in which different surrogate models are necessary to capture the relevant model behavior and ensure computational tractability: ignition-point inversion (approximated with deep flow maps) and smoke impacts (approximated with a basis of polynomial operators). The implementation of both surrogate models is available in open-source software.

Zachary Morrow
Sandia National Laboratories
zbmorro@sandia.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Joseph Hart, Bart van Bloemen Wanders, Daniel Krofcheck
Sandia National Laboratories
joshart@sandia.gov, bartv@sandia.gov, djkrofc@sandia.gov

MS77

Optimal Experimental Design for Digital Twins Using a Network-Flow Approach

Efficient data assimilation is a significant bottleneck towards the development of highly scalable and reliable digital twins. Uncertainty quantification and optimization becomes particularly challenging when either the nature of the system or cost limitations prevent the deployment of a large number of sensors for data collection. We propose a novel trajectory design approach for the optimal experimental design (OED) of contiguous data acquisition based on a network-flow formulation. In contrast to conventional control strategies, we discretize the spatial domain into a set of network nodes and identify the optimal path between nodes to minimize uncertainty. This approach allows for a more thorough exploration of the design space than what multi-start local optimizations can do. The network-flow approach we propose translates into a challenging discrete optimization problem, whose objective is expensive to evaluate, highly nonlinear, and can be non-smooth. Moreover, the objective function needs to be evaluated over the entire path instead of being summed edge-by-edge. With an increasing number of edges, the number of possible paths rapidly increases and the combinatorial problem becomes intractable. We showcase a branch-and-bound algorithm we developed, which leverages tailored bounds and constraints to tackle the optimization at a tractable computational cost.

Marco Mangano
Argonne National Laboratory
mmangano@anl.gov

Paul Manns
Mathematik
TU Dortmund
paul.manns@tu-dortmund.de

Sven Leyffer
Argonne National Laboratory
leyffer@anl.gov

Sven Leyffer, Carlo Graziani, Rebecca Durst
Argonne National Laboratory
leyffer@anl.gov, cgraziani@anl.gov, rdurst@anl.gov

MS77

Adaptive Sensor Optimization with Performance Guarantees for Digital Twins

The scalable optimization of sparse sensor placement is crucial for prediction, sensing and control of safety-critical engineering systems. Identifying the fewest possible sensor measurements often translates into sparsity penalties on optimization design variables. Current approaches rely on relaxation, heuristics, and trial-and-error selection of hyperparameters, which challenge the interpretability and verification of resulting sensor placements. In this talk, we propose a statistical mechanics approach for robust sparse sensor placement, using free energies to analyze the optimization landscape, optimize hyperparameters and quantify uncertainty with respect to noise and data volume. We illustrate how this approach provides optimization landscapes and critical noise regimes for reconstruction of high-dimensional fields from sparse sensors and data-driven pri-

ors. We showcase how these tools can be used for constrained optimization of sensor placement in nuclear energy digital twins.

Krithika Manohar
University of Washington
kmanohar@uw.edu

MS77

Risk-aware Nonlinear Bayesian Experimental Design

Bayesian optimal experimental design typically aims to maximize expected information gain, equivalent to maximizing mutual information. This reflects a risk-neutral perspective, treating all possible outcomes based on their average Kullback–Leibler (KL) divergence between prior and posterior. However, in settings where risk aversion is important, expectation may not be the most appropriate objective. We propose a new criterion based on superquantile information gain, which captures the tail behavior of the information distribution and aligns with risk-averse decision-making. We also develop a Monte Carlo estimator tailored for this criterion. The approach is demonstrated on synthetic benchmarks and a sensor placement problem in fluid dynamics.

Ellen Meyer, Xun Huan
University of Michigan
ecmeyer@umich.edu, xhuan@umich.edu

MS78

Sample-Based Uncertainty Quantification for Inverse Problems with Poisson Data

Inverse problems arise when physics-based models are fit to observational data. The unknown parameters in such models are often spatially distributed and become large-scale after discretization. The resulting parameter estimation problems are typically ill-posed, meaning that solutions are unstable with respect to errors in the data. In this talk, I will focus on inverse problems where the data follows a Poisson distribution, a common scenario in imaging applications. I will formulate the Poisson inverse problem within a Bayesian inference framework, allowing for uncertainty quantification through the posterior distribution. I will then present MCMC methods for sampling from the posterior in the Poisson case that incorporate nonnegativity constraints, which are natural in imaging applications.

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

MS78

Efficient Parameter Estimation for Separable Nonlinear Inverse Problems

Good estimates of hyperparameters can be critical for solving inverse problems accurately, but estimating these hyperparameters often requires solving a challenging and expensive optimization problem involving log determinants. In this work, we consider two optimization approaches, both of which involve stochastic average approximations (SAA) of the objective function. The first approach uses a Monte Carlo approach to approximate the objective function and gradient. The second approach is a majorization-minimization approach that replaces a challenging optimization problem with a sequence of simpler ones. Specifically, we use a majorization function for the log-determinant term and combine it with a randomized trace estimator to form a majorant. Both of these methods can be used within a general framework for hyperparameter estimation that includes both estimating standard hyperparameters (e.g., noise and prior variances) and estimating model parameters (e.g., unknown parameters that define the forward process). That is, we consider separable nonlinear inverse problems, where the forward model assumes a parametric form, and the goal is to improve the model parameters to more accurately solve the resulting linear inverse problem. By treating the model parameters as hyperparameters and using the proposed approaches for hyperparameter estimation, we provide an efficient approach to estimate model parameters in separable nonlinear inverse problems.

Elle Buser
Emory University
ebuser@emory.edu

Julianne Chung
Department of Mathematics, Emory University, USA
julianne.mei-lynn.chung@emory.edu

Arvind Saibaba, Hugo Diaz-Norambuena
North Carolina State University
asaibab@ncsu.edu, hsdiazno@ncsu.edu

MS78

Spotlight Inversion by Orthogonal Projections in a Bayesian Framework

In inverse problems, the goal is to estimate unknown parameters from indirect noisy observations. It is not uncommon that the forward model assigning the observed variables to given values of the unknowns depend on variables that are not of primary interest, often referred to as nuisance parameters. In this article, we consider linear inverse problems, and propose a novel technique based on orthogonal projections to eliminate, or at least mitigate, the contribution of the nuisance parameters on the data. The approach is referred to as spotlight inversion, as it allows to focus on the part of primary interest of the unknown parameter, leaving the uninteresting part in the shadow. Focusing the spotlight on the portion of the solution of interest may lead to very substantial reductions in the computational complexity, as illustrated with computed examples.

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

Nuutti Hyvonen
Aalto University
nuutti.hyvonen@aalto.fi

Ville P. Kolehmainen
Department of Technical Physics
University of Eastern Finland
ville.kolehmainen@uef.fi

Erkki Somersalo
Case Western Reserve University

ejs49@case.edu

MS78

Variable Projection through a Bayesian Lens

Many inverse problems, including convolution-based models arising in imaging, involve operators whose kernels depend on unknown parameters. In such settings, inference requires estimating these parameters while accounting for uncertainty. When the forward model is linear in some variables and nonlinear in others, the resulting posterior distribution exhibits exploitable structure. In this presentation, we show that the classical variable projection method admits a natural Bayesian interpretation. By marginalizing over the linear variables, we obtain a reduced posterior over the nonlinear parameters that coincides with the variable projection objective. This perspective enables efficient sampling in the reduced space, improving computational efficiency while retaining principled uncertainty quantification for inverse problems with separable structure.

Jordan T. Dworaczyk, Malena I. Espanol
Arizona State University
jtdworac@asu.edu, malena.espanol@asu.edu

MS79

Reconstructing Cardiac Dynamics and Model Parameters Using Data Assimilation

Cardiac ventricular action potentials arise from nonlinear interactions between transmembrane voltage, ionic currents, and intracellular ion concentrations. Current experimental methods cannot simultaneously record membrane potential and multiple currents, limiting understanding of how perturbations such as pacing changes or drug treatments initiate or prevent arrhythmias. Biophysical models simulate these dynamics but rely on averaged parameters, reducing their ability to capture biological heterogeneity. We address this using data assimilation, a computational approach that integrates models with noisy observations to jointly estimate state variables and parameters. A heterogeneous virtual population of guinea pig ventricular cells was created by varying ionic conductances, with Gaussian noise added to simulated voltage and calcium signals. The method reconstructed the full state space with error below noise levels and accurately estimated conductances. Compared with two post hoc optimization methods, data assimilation achieved higher accuracy with less computation. Applied to a virtual population of human iPSC-derived ventricular cells, it successfully reconstructed complex action potential morphologies and key parameters. These results show that data assimilation enables precise state reconstruction and parameter estimation in heterogeneous populations, enhancing the predictive power of biophysical models for studying arrhythmia mechanisms.

Mario Mendez Perez
College of Engineering
The Ohio State University
mendezperez.1@osu.edu

Seth H. Weinberg
Ohio State University
weinberg.147@osu.edu

MS79

Multilevel MCMC for Bayesian Inference of IART Scar Parameters with Uncertainty

Cardiac arrhythmia arise from abnormal electrical activity in the heart and are associated with a large number of deaths. Their underlying pathophysiology varies between individuals, necessitating patient-specific treatment. Personalized models or cardiac digital twins offer a promising new approach. However, current imaging techniques limit the fidelity with which individual cardiac anatomy and function can be captured, making uncertainty quantification (UQ) in model predictions essential. We employ Bayesian inference for parameter estimation in a cardiac electrophysiology model and for UQ that accounts for measurement and discretization errors. In a two-dimensional cardiac tissue slab, we simulate the electrical signal propagation that coordinates heart contraction. From mapping catheter data, we aim to estimate the shape parameters of electrically inactive regions, or scars, which disrupt signal propagation and can lead to arrhythmia. In our talk, we propose a Markov chain Monte Carlo method that requires a minimal number of samples for posterior estimation by defining a likelihood based on summary statistics and by using an adaptive proposal distribution. To account for the discretization error, we inflate the likelihood variance, and we tailor the meshing strategies used for the forward simulations in the accept-reject step. Further computational efficiency is achieved through a multilevel method that substitutes expensive simulations with cheaper approximations.

Maarten Volkaerts
Katholieke Universiteit Leuven
maarten.volkaerts@kuleuven.be

Marie Cloet
Department of Computer Science
KU Leuven
marie.cloet@kuleuven.be

Piet Claus
Cardiovascular Imaging and Dynamics
KU Leuven
piet.claus@kuleuven.be

Hans Dierckx
Leids Universitair Medisch Centrum
h.j.f.dierckx@lumc.nl

Giovanni Samaey
KU Leuven
giovanni.samaey@kuleuven.be

MS79

Active Learning for Uncertainty Quantification in Personalized Cardiac Models

Personalized computer simulation models provide physicians with a patient-specific virtual organ to predict therapeutic responses, at minimal to no risk to the patient. However, clinical translation of these models has been slow. A key barrier is the variability in model predictions that, if not quantified, will hamper decision support. Probabilistic inference in this context is challenging, because standard Markov Chain Monte Carlo (MCMC) sampling requires repeated model simulations that are computationally infeasible. Surrogate modeling can be an effective solution by building an efficient approximation of the expensive simulation model. To build a globally accurate approximation of a highly nonlinear model, however, is computa-

tionally infeasible without sacrificing local approximation accuracy. In this talk, we describe our effort in cast surrogate modeling as a machine-learning problem, in which the training data are virtually infinite (in a continuous space) but expensive to obtain (incurring model evaluations). We describe the concept of active surrogate modeling where surrogate construction will actively search the optimal training points at which to query the model. This high-dimensional active search is further embedded into a low-dimensional latent space through probabilistic generative modeling, which enables surrogate-accelerated MCMC sampling. We demonstrate the efficacy of these developments in the context of patient-specific cardiac modeling.

Linwei Wang
Rochester Institute of Technology
linwei.wang@rit.edu

MS80

Approximate Bayesian Inference for Cislunar Tracking Using Alternate Coordinate Representations

The cislunar dynamical environment is characterized by highly nonlinear and chaotic behavior. Conventional approaches to filtering and tracking, which tend to impose simplifying assumptions on the dynamics and uncertainty realization, are no longer appropriate. The Modified Generalized Equinoctial Orbital Elements (M-GEqOEs) are a promising technique for high-fidelity propagation of states and uncertainty in cislunar space. These orbital elements may be embedded with any necessary perturbations, conservative or non-conservative, and the element set is defined for the entire cislunar domain. In addition, uncertainty evolved using the M-GEqOEs preserves Gaussian behavior for longer propagation times than conventional Cartesian propagation, yielding improvements in uncertainty realism. The downside of the M-GEqOEs (or other non-Cartesian coordinates) is that they tend to introduce additional nonlinearities relative to conventional measurement sets, such as range and range-rate measurements or line-of-sight measurements. The added nonlinearity tends to deteriorate the performance of standard filtering and tracking approaches. This work combines the M-GEqOE methodology with recent advances in approximate Bayesian filtering to formulate a parameter flow approach that ameliorates the issues of coordinate system selection in the measurement update. Simulation results are considered for standard cislunar test cases to demonstrate the efficacy of the developed approach.

Kyle DeMars
Texas A&M
kdemars@purdue.edu

Maaninee Gupta
Texas A&M University
maanineegupta@tamu.edu

MS80

Advances in Kernel-Based Ensemble Gaussian Mixture Filtering with Applications in Astrodynamics

This work advances techniques for onboard nonlinear filtering with applications in cislunar astrodynamics. An effective onboard filter must be accurate, consistent, efficient, and verifiable. Recent work has shown that the kernel-based Ensemble Gaussian Mixture Filter (EnGMF) is an accurate and consistent nonlinear filter capable of handling non-Gaussian probability distributions. However, its fixed number of Gaussian components restricts efficiency and its random resampling step prevents verifiability. To address these issues, the EnGMF has recently been modified such that it adaptively controls its number of Gaussian components (improving efficiency) and it deterministically resamples components in its resampling step (ensuring verifiability). The presented applications revolve around cislunar orbits, demonstrating the need for nonlinear filtering methods and their potential for deep-space operations. Simulation results indicate that the presented advancements of the EnGMF not only maintain the accuracy and consistency of the traditional EnGMF, but also improve its efficiency and ensure its verifiability—a significant step forward in developing onboard nonlinear filters with applications in cislunar astrodynamics.

Dalton Durant
University of Texas at Austin
ddurant@utexas.edu

Felipe Giraldo Grueso
The University of Texas at Austin
fgiraldo@utexas.edu

Renato Zanetti
University of Texas at Austin
Department of Aerospace Engineering and Engineering Mechanic
renato@utexas.edu

Andrey A. Popov
University of Hawai'i at Manoa, U.S.
apopov@hawaii.edu

Uwe Hanebeck
Karlsruhe Institute of Technology
uwe.hanebeck@kit.edu

MS80

Polynomial Updates for the Unscented Kalman Filter

Common filters are usually based on the linear approximation of the optimal minimum mean square error estimator. The Unscented Kalman Filter handles nonlinearity of the dynamics and of the measurement equation through the unscented transformation, but remains a linear estimator, meaning that the state estimate is a linear function of the measurement. This talk proposes a polynomial approximation of the optimal estimator, improving the update step of the common UKF. The new update retains the structure of the linear counterpart but includes information from the measurement powers to obtain a more accurate estimate. The result is a polynomial update that better approximates the true Minimum Mean Square Error (MMSE), drastically improving the accuracy of the filter, and that can be generalized to any arbitrary order. Whenever needed, the prediction of uncertainties in the UKF is enhanced with the conjugate unscented transformation (CUT), to best obtain the high-order central moments predictions required to perform a polynomial update. Numerical results show the benefits in accuracy and robustness of the new technique, which can be generalized to upgrade other linear estimators to their polynomial versions and to work with

non-Gaussian noises.

Simone Servadio, Chiran Cherian
Iowa State University
servadio@iastate.edu, cbckbc@iastate.edu

MS80

Polynomial Chaos-Based Estimation Leveraging Dimension Reduction and the Active Subspace

With increasing mission complexity and space activity prompting the desire for more capable orbit determination methods, uncertainty quantification in astrodynamics has increasingly relied on non-Gaussian estimators and surrogate modeling techniques. One type of surrogate model that has seen increasing usage in this field is polynomial chaos, a response surface method that maps a set of input uncertainties to an output space using orthogonal polynomial basis functions. Here, we present the Intrusive Polynomial Chaos Filter (IPCF), a nonlinear estimator utilizing polynomial chaos expansions (PCEs) to represent the quantities of interest. Like many complex uncertainty quantification methods, PCEs are vulnerable to the curse of dimensionality, where growth in dimension and degree disproportionately increase complexity. In the IPCFs case, the nonlinear update increases the solutions degree, and new random inputs are added each time the surrogate is propagated or updated with a measurement. To control the surrogates growth, the degree and dimension are reduced to those of the original prior PCE after each update. In this work, we present two methods of input dimension reduction, each with different advantages. The first leverages sensitivity analysis via the Sobol indices to reduce the surrogate cost-effectively. The second identifies and exploits the PCEs active subspace to preserve more information at a higher computational cost.

Zee Toler
The University of Texas at Austin
zee.toler@utexas.edu

MS81

A Natural Primal-Dual Hybrid Gradient Method for Adversarial Neural Network Training on Solving Partial Differential Equations

We propose a scalable deep learningbased algorithm for solving PDEs that integrates preconditioned gradient methods with adversarial training. By multiplying the PDE with a test function, we reformulate the problem as an infsup saddle-point problem and apply the PrimalDual Hybrid Gradient (PDHG) algorithm. To enhance efficiency and stability, we introduce appropriate preconditioning operators into the proximal steps of the PDHG algorithm, yielding a natural gradient ascentdescent scheme for updating neural network parameters. The resulting natural gradients are evaluated efficiently using the Krylov subspace method (MINRES) involving iterative matrixvector products. Convergence guarantees are established for time-continuous algorithms when applied to linear equations, particularly elliptic PDEs in divergence form. We validate the method on a wide range of linear and nonlinear PDEs with dimensions ranging from 1 to 50. Numerical results demonstrate more stable convergence and higher accuracy of our approach than several widely used deep learning-based solvers.

Shu Liu
University of California, Los Angeles, U.S.
sl25bn@fsu.edu

MS81

Transformers Learn Variable-Order Markov Chains In-Context

Large language models show strong in-context learning (ICL), but how transformers implement it in complex settings is unclear. Prior work targets fixed-order Markov chains (FOMC); natural language fits variable-order Markov chains (VOMC), or context trees (CTs). We study VOMC ICL as data compression on small alphabets and low orders, using context-tree weighting (CTW) and prediction by partial matching (PPM) as baselines; CTW is Bayesian-optimal under a class of CTW priors. Empirically, (1) transformers can compress VOMC in context, while PPM degrades; (2) depth matters littleeven two layers work well; and (3) models trained/tested on non-CTW priors can beat CTW. To explain this, we inspect attention and extract two mechanisms, yielding two constructions: (1) a D+2-layer transformer that accurately mimics CTW for CTs of maximum order D; and (2) a 2-layer transformer that uses the feed-forward network for probability blending. Unlike the FOMC setting, an explicit counting mechanism is key. Implemented synthetic layers produce hybrid transformers that match or exceed standard transformers VOMC ICL while using far fewer parameters.

Chao Tian
Texas A&M University
chao.tian@ece.tamu.edu

MS81

Active Subspace Inference Dynamics in Bayesian Neural Networks

TBD

Nathan Urban
Brookhaven National Laboratory
nurban@bnl.gov

MS81

In-Context Learning of Linear Dynamical Systems with Transformers: Error Bounds and Depth-Separation

This paper investigates approximation-theoretic aspects of the in-context learning capability of the transformers in representing a family of noisy linear dynamical systems. Our first theoretical result establishes an upper bound on the approximation error of multi-layer transformers with respect to an -testing loss uniformly defined across tasks. This result demonstrates that transformers with logarithmic depth can achieve error bounds comparable with those of the least-squares estimator. In contrast, our second result establishes a non-diminishing lower bound on the approximation error for a class of single-layer linear transformers, which suggests a depth-separation phenomenon for transformers in the in-context learning of dynamical systems. Moreover, this second result uncovers a critical distinction in the approximation power of single-layer linear transformers when learning from IID versus non-IID data.

Yuxuan Zhao
University of Minnesota Twin Cities

zhao1895@umn.edu

MS82

Pair Approximation for Continuous-Valued Dynamics on Networks

Dynamics on networks arise in diverse contexts such as information spreading, disease transmission, epidemiology, opinion dynamics, and statistical physics. In these systems, agents (nodes) interact with their neighbors through network edges, and their states evolve in time depending on the states of their neighbors. The network structure, therefore, plays an essential role in shaping the overall dynamics. When networks are large or highly heterogeneous, direct computations become infeasible and mean-field approximations often fail to capture essential correlations. Pair approximation offers a powerful reduction framework by incorporating network heterogeneity through degree-based correlations, but most existing analyses focus on systems with discrete-valued states, where agents take values from a finite set. In this work, we extend the framework to continuous-valued dynamics. We introduce a family of degree-based density functions to characterize the state distributions of nodes with the same degree and derive the corresponding pair approximation for these density functions.

Weiqi Chu
University of Massachusetts Amherst
weiqichu@umass.edu

MS82

Learning Solution Operator of Dynamical Systems with Diffusion Maps Kernel Ridge Regression

In this talk, I will discuss a simple kernel ridge regression (KRR) framework with a dynamic-aware validation strategy for long-term prediction of complex dynamical systems. The proposed approach implicitly adapts to the intrinsic geometry of the system's invariant set, without requiring explicit manifold reconstruction or attractor modeling, procedures that often limit predictive performance. I will demonstrate the robustness of this approach across a broad range of systems, including smooth manifolds, chaotic attractors, and high-dimensional spatiotemporal flows, DM-KRR consistently outperforms state-of-the-art random feature, neural-network and operator-learning methods in both accuracy and data efficiency.

John Harlim
Pennsylvania State University
jharlim@psu.edu

MS82

On the Dimensionality of Dynamics on Complex Networks: from Exact Reductions to Approximate Ones

Behind the blur of high-dimensional nonlinear dynamics and the intricate organization of complex systems lie essential mechanisms that drive the emergence of macroscopic phenomena. Diverse approaches to dimensionality reduction have been developed to uncover these mechanisms. Reducing the dimensionality of dynamical systems on complex networks is challenging, though exact or approximate reductions are sometimes possible. This talk presents concrete examples and their implications. First, we consider the Kuramoto model with heterogeneous phase lags on general weighted, directed, signed graph. Using Koopman theory, we derive necessary and sufficient conditions for the existence of distinct constants of motion in the model, including new ones generated from Lie symmetries. These results reveal a broad class of network motifs that support conserved quantities, thereby enabling exact dimension reductions and providing new insights into synchronization. Second, we show that complex systems are often modeled using dynamics that depend on low-rank matrices describing the networks of interactions. We present three indicators of the low-rank hypothesis and show its ubiquity in random network models. We then verify the hypothesis for empirical networks of various origins and examine its repercussions for nonlinear dynamics. In particular, we demonstrate that networks described by low rank matrices cause the dynamics they support to collapse onto a much lower-dimensional manifold.

Vincent Thibeault
Université Laval
vincent.thibeault.1@ulaval.ca

MS83

Operator Learning Approaches for Inverse Problems in Imaging

Inverse problems in imaging often involve reconstructing unknown quantities from indirect and noisy measurements, and their solution is typically hindered by expensive forward computations and ill-posedness. Recent advances in operator learning provide new perspectives on how to address these challenges. In this talk, we will discuss how neural operator architectures can be used both to accelerate classical Bayesian inference methods and to learn novel operator-to-function mappings arising in imaging applications. In the talk, we will discuss some theoretical approximation guarantees and also illustrate these ideas through examples from medical imaging.

Anuj Abhishek
Case Western University
axa1828@case.edu

MS83

Reduced Modeling of Chaotic, Turbulent, and Stochastic Systems via Population Dynamics

Learning models of time-dependent processes that generalize across initial conditions and parameter regimes is a key challenge in machine learning and the computational sciences. For chaotic, turbulent, and stochastic systems, modeling the dynamics of individual trajectories can be exceedingly challenging because trajectories can be erratic and irregular, and in stochastic settings may even be nowhere differentiable. Instead, we propose learning population dynamics, which model how the distribution of the system states evolves over time. By learning population dynamics, we deliberately discard trajectory-specific information to obtain dynamics that are smoother and more well behaved than the underlying sample trajectories. We argue that this loss of information is often an acceptable, principled form of reduced modeling because population-level statistics are maintained, and these are often the quantities that are of interest in scientific and engineering settings.

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University

pehersto@cims.nyu.edu

MS83

Scalable Nonlinear Manifold Reduced Order Models: Domain Decomposition and Cross-PDE Expressivity

The domain decomposition (DD) nonlinear-manifold reduced-order model (NM-ROM) represents a computationally efficient strategy that integrates underlying physics principles into a neural networkbased, data- driven framework. Compared to linear subspace methods, NM-ROMs pro- vide superior expressivity and enhanced reconstruction capabilities, while the DD strategy enables cost-effective, parallel training of autoencoders by partitioning the domain into algebraic subdomains. In this work, we develop and assess DD-NM-ROMs for the two-dimensional time-dependent Burgers equation and the diffusion equation, and compare them with their linear subspace counterparts (DD-LS-ROMs). Both approaches em- ploy a bottom-up strategy: ROMs are trained on smaller subdomains and subsequently deployed on larger, composable domains. Remarkably, the neural network manifold representation demonstrates the ability to generalize across PDEs: a single network trained jointly on Burgers and diffusion data can accurately predict unseen solutions of each equation individually. Numerical experiments show that DD-NM-ROMs consistently outperform DD-LS-ROMs in both accuracy and computational speed-up, while also enabling stable extrapolation to larger spatial domains. These findings highlight the potential of DD-NM-ROMs as a scalable and expressive framework for reduced-order modeling of diverse PDE systems.

Ping-Hsuan Tsai
Virginia Tech
pinghsuan@vt.edu

MS83

Generative Ensembles for UQ: A Diffusion Model Approach to Forward and Inverse Problems

Score-based diffusion models offer a unified, probabilistic framework for solving forward and inverse problems in scientific computing (Jacobsen, Zhuang, Duraisamy; SISC, 2025). This presentation demonstrates their power for robust uncertainty quantification (UQ) through contributions in field reconstruction and large- scale weather forecasting. For inverse problems, our spatially-aware diffusion model reconstructs global fields from sparse data by generating an ensemble of plausible realizations (Zhuang, Cheng, Duraisamy; CMAME, 2025). We demonstrate that the ensemble mean converges to the deterministic solution, providing a well-behaved framework for UQ and covariance estimation for data assimilation. For forward problems, we showcase our LaDCast framework, the first latent diffusion model for global weather forecasting (Zhuang, Cheng, Duraisamy; 2025). It produces a 10-day forecast in under 10 seconds on a single GPU and generates an entire forecast ensemble without artificial noise injections. This provides a computationally efficient and low-bias approach to UQ for a grand-challenge problem. These examples highlight a pathway toward diffusion-based probabilistic foundation models that are scalable and inherently suited for the UQ demands of modern scientific computing.

Yilin Zhuang
University of Michigan
ylzhuang@umich.edu

MS85

Conditional Generation of Rare Events with Diffusion Models

Diffusion models have emerged as powerful tools to generate scenarios design and optimization, offering the ability to generate diverse (high recall) and high-quality (high precision) samples. For critical energy infrastructures, design constraints are often framed as resilience to rare scenarios, typically expressed as probabilistic constraints. This motivates the need for probability-conditioned scenario generation. In this work, the problem of generating turbulent fluid flow scenarios conditioned on their probability is explored. This task relies on two separate steps: a) labeling data snapshots with a probability b) conditioning the generation on said probability. For step a), two approaches are proposed: one that relies on normalizing flows to label the data samples; and one that labels data samples from the diffusion process via a score-based likelihood reconstruction. It is shown that while the rarity of the generated scenarios can be controlled, generating events more rare than the ones observed in the dataset is not possible. In view of this result, a multifidelity training approach for the diffusion model is proposed and discussed.

Malik Hassanaly
National Renewable Energy Laboratory
malik.hassanaly@nrel.gov

Graham Harper
Computational Mathematics Dept.
Sandia National Laboratories
gbharpe@sandia.gov

Zach Grey
National Institute of Standards and Technology
Applied and Computational Mathematics Division
zachary.grey@nist.gov

Demi Qin, Alex Rybchuk
National Renewable Energy Laboratory
demi.qin@nrel.gov, alex.rybchuk@nrel.gov

Cristian J. Villatoro
University of Notre Dame
cvillat2@nd.edu

Andrew Glaws
National Renewable Energy Laboratory
15013 Denver West Parkway Golden, CO 80401-3305
andrew.glaws@nrel.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Ryan King
NREL
ryan.king@nrel.gov

MS85

A Training-Free Conditional Diffusion Model for Learning Stochastic Dynamical Systems

This study introduces a training-free conditional diffusion model for learning unknown stochastic differential equa-

tions (SDEs) using data. The proposed approach addresses key challenges in computational efficiency and accuracy for modeling SDEs by utilizing a score-based diffusion model to approximate their stochastic flow map. Unlike the existing methods, this technique is based on an analytically derived closed-form exact score function, which can be efficiently estimated by a Monte Carlo method using trajectory data, and hence eliminates the need for neural network training to learn the score function. By generating labeled data through solving the corresponding reverse ordinary differential equation, the approach enables supervised learning of the flow map. Extensive numerical experiments across various SDE types, including linear, nonlinear, and multi-dimensional systems, demonstrate the versatility and effectiveness of the method. The learned models exhibit significant improvements in predicting both short-term and long-term behaviors of unknown stochastic systems, often surpassing baseline methods like GANs in estimating drift and diffusion coefficients.

Yanfang Liu
Middle Tennessee State University
yanfang.liu@mtsu.edu

Yuan Chen, Dongbin Xiu
Ohio State University
chen.11050@buckeyemail.osu.edu, xiu.16@osu.edu

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS85

Multimodal Data Assimilation Using Score-Based Diffusion Models

Score-based diffusion modeling is a generative machine learning algorithm that can be used to sample from complex distributions. They achieve this by learning a score function, i.e., the gradient of the log-probability density of the data, and reversing a noising process using the same. Once trained, score-based diffusion models not only generate new samples but also enable zero-shot conditioning of the generated samples on observed data. This promises a novel paradigm for data and model fusion, wherein the implicitly learned distributions of pretrained score-based diffusion models can be updated given the availability of online data in a Bayesian formulation. In this article, we apply such a concept to the super-resolution of a high-dimensional dynamical system, given the real-time availability of low-resolution and experimentally observed sparse sensor measurements from multimodal data. Additional analysis on how score-based sampling can be used for uncertainty estimates is also provided. Our experiments are performed for a super-resolution task that generates the ERA5 atmospheric dataset given sparse observations from a coarse-grained representation of the same and/or from unstructured experimental observations of the IGRA radiosonde dataset. We also demonstrate accurate recovery of the high dimensional state given multiple sources of low-fidelity measurements and AI-based numerical weather prediction models.

Romit Maulik, Dibyajyoti Chakraborty
Pennsylvania State University
rmm7011@psu.edu, d.chakraborty@psu.edu

Haiwen Guan
Penn State University
hzg18@psu.edu

MS85

Constraint-Guided Conditional Diffusion for Power-Grid Generation

Maps of distribution grids are often incomplete, proprietary, or out of date, yet cities still need to plan new neighborhoods, move lines underground, integrate local energy technologies (rooftop solar, home batteries, EV charging), and rebuild after storms. We frame grid synthesis as conditional generative modeling: a diffusion model that learns a distribution over feasible feeder topologies and samples many plausible designs. Sampling is guided by grid constraintsradial structure, equipment/connection limits, and basic voltage/thermal checksso the generated networks remain physically plausible while explicitly capturing uncertainty. The result is a practical scenario engine: instead of one best network, we produce ensembles ready for standard planning workflows (power flow, hosting-capacity, and resilience studies). For evaluation, we use the generated graphs as inputs to a pipeline that constructs complete, simulation-ready synthetic distribution feeders from map and context data, showing improved realism, diversity, and downstream utility over prior GNN-based topology generators. To our knowledge, this is the first conditional diffusion approach for grid topology generation, offering a data-light, uncertainty-aware path to modern grid planning.

Demi Qin
National Renewable Energy Laboratory
demi.qin@nrel.gov

Andrew Glaws
National Renewable Energy Laboratory
15013 Denver West Parkway Golden, CO 80401-3305
andrew.glaws@nrel.gov

Ryan King
NREL
ryan.king@nrel.gov

Aadil Latif
National Renewable Energy Laboratory
aadil.latif@nrel.gov

MS86

Towards the Construction of Inverse Triangular Transport Maps from Multifidelity Information

Transport maps provide deterministic couplings between probability distributions and are broadly useful in uncertainty quantification tasks. Triangular transport maps, distinguished by their eponymous triangular structure, possess several desirable numerical and theoretical properties. Map construction can be done using samples from the target distribution in a process called inverse transport, but the accuracy of these maps is often contingent on having many high-fidelity training samples, which can be costly to obtain. To address this, we build maps from multifidelity information, leveraging an ensemble of information sources of varying cost and accuracy. We propose two multifidelity inverse transport methods. The first is a hierarchical approach based on deep transport maps and inspired by [1]. We train maps to move between adjacent distributions within an assumed fidelity hierarchy. The second method trains a single map that incorporates all information sources simultaneously through multifidelity correc-

tions to its parameterization. We compare both approaches across various problems, demonstrating their effectiveness in reducing high-fidelity data requirements in inverse transport. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525 [1] Benjamin Peherstorfer and Youssef Marzouk. A transport-based multifidelity preconditioner for Markov chain Monte Carlo. In: Advances in Computational Mathematics 45 (2019), pp. 23212348.

Owen N. Davis
University of New Mexico
ondavis@sandia.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

MS86

Estimating Probability Distributions Using Multi-Fidelity Methods

Multi-fidelity sampling methods are efficient strategies to perform forward uncertainty quantification in a variety of applications. These methods are based on the classical Monte Carlo estimator, but fuse predictions from two or more models of varying fidelity to obtain estimators with a greater precision given a prescribed computational budget. Despite their success, most of these multi-fidelity methods are designed to estimate statistical moments, e.g., mean and variance, while few previous works have focused on estimating probability distributions or quantiles. In this work, we extend the approximate control variates (ACV) approach to estimate cumulative density functions using multi-output estimation. A novel interpolation method for smoothing out statistical noise is also introduced to further improve estimator accuracy. We present results on the evaluation of distributions for additively manufactured Ni-base superalloys.

Thomas Dixon
University of Michigan
tdixono@umich.edu

Geoffrey Bomarito
NASA Langley Research Center
geoffrey.f.bomarito@nasa.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

James Warner
Nasa Langley Research Center
james.e.warner@nasa.gov

Joshua Pribe
Analytical Mechanics Associates
joshua.pribe@nasa.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

Alex Gorodetsky
University of Michigan
goroda@umich.edu

MS86

Multi-Fidelity Transport-Map Augmented MCMC

Model calibration in computational physics is often hindered by the prohibitive cost of high-fidelity simulations. Traditional approaches mitigate this by replacing expensive solvers with surrogate models, but this substitution can reduce robustness and incur bias. We propose a multi-fidelity Bayesian inference framework that fuses information from models of varying resolution by combining Markov chain Monte Carlo (MCMC) with multilevel Monte Carlo (MLMC) estimators. The key innovation is a transport map-based coupling in a shared reference space, which induces strong correlations among posterior samples across fidelity levels, leading to accurate and low-variance estimators. This construction avoids the need for surrogate replacement, while ensuring that all models contribute coherently to the inference in a physically consistent manner. We demonstrate our approach in turbulence model calibration, where Reynolds-averaged NavierStokes (RANS) models require inference of uncertain closure coefficients to improve predictive reliability in nonlinear flow regimes. By leveraging a surrogate and a RANS solver, the framework delivers calibrated posteriors and predictive uncertainty bounds at substantially reduced computational cost compared to single-fidelity methods. These results demonstrate that multi-fidelity calibration achieves significant speedups over traditional Bayesian approaches, making large-scale uncertainty quantification in complex models more practical.

Sanjan Muchandimath, Joaquim Martins, Alex Gorodetsky
University of Michigan
sanjancm@umich.edu, jrram@umih.edu, goroda@umich.edu

MS86

Uncertainty-aware Mixture of Experts Surrogate Models for Computational Material Science

We propose a framework for uncertainty aware mixture-of-experts surrogate constitutive models with a focus on computational material science. This framework allows for the integration of multiple heterogeneous models, in a multifidelity framework. Each model addresses specific aspects of the material behavior and are combined in a multi-fidelity framework for a more accurate representation of the material response to different physical regimes. We endow this framework with uncertainty via Langevin sampling by introducing a stochastic differential equation with a trainable drift function corresponding to the log-posterior distribution of the surrogate model parameters. We will discuss computational expense tradeoffs between the dimensionality of the surrogate constitutive model and the associated posterior distribution model, and the number of samples generated by the Langevin process, and their impact on model predictions for practical computational material science exemplars.

Cosmin Safta, Pieterjan Robbe, Reese Jones
Sandia National Laboratories
csafta@sandia.gov, pmrobbe@sandia.gov, rjones@sandia.gov

Christophe Bonneville
Sandia National Laboratories, U.S.

cpbonne@sandia.gov

Habib N Najm
Sandia National Labs
USA
hnnajm@sandia.gov

Laurent Capolungo
Los Alamos National Laboratory
laurent@lanl.gov

MS87

Efficient Reliability Analysis of Stochastic Simulators via Active Learning of Heteroscedastic GP

Reliability analysis of computationally expensive stochastic simulators is challenged by the need to disentangle epistemic (modeling) uncertainty from aleatoric (intrinsic noise) uncertainty. Standard surrogate models often fail this task, especially under tight computational budgets where simulation replications are infeasible. We introduce an active learning framework, AL-NR-HGP, that efficiently addresses this problem. Our approach is founded on a non-replicated heteroscedastic Gaussian process (NR-HGP), which uses coupled GPs to simultaneously learn a system's mean response and its input-dependent noise from a single set of non-replicated data. The learning process is driven by a novel learning function derived from a first-order sensitivity analysis of the conditional failure probability. This function intelligently guides exploration by decomposing the posterior variance into its epistemic and aleatoric contributions, prioritizing points that most effectively reduce uncertainty about the limit-state boundary. For parallel execution, a weighted k-means batching strategy ensures the selection of diverse and informative points. On benchmark problems, AL-NR-HGP demonstrates superior accuracy and efficiency, providing more robust failure probability estimates than crude Monte Carlo and competing surrogate-based methods.

Han Peng, Jize Zhang
Hong Kong University of Science and Technology
hpengak@connect.ust.hk, cejize@ust.hk

MS87

Deep Learning Frameworks for Metamodeling of Nonlinear Stochastic Dynamic Systems under Parametric and Predictive Uncertainty

Metamodeling of high-dimensional, nonlinear dynamic structural systems subject to natural hazards presents formidable challenges, especially when simultaneously accounting for uncertainties in the external dynamic loads and structural parameters. While uncertainties related to external dynamic loads from natural hazards have been successfully incorporated, few formulations have simultaneously considered loading and parameter uncertainties. To address this gap, three metamodeling frameworks are formulated in this work, each coupling a feature-extraction moduleimplemented through a multilayer perceptron (MLP), a message-passing neural network (MPNN), or an autoencoder (AE)with a long short-term memory (LSTM) network using Monte Carlo dropout and a negative log-likelihood loss. The resulting architectures (MLP-LSTM, MPNN-LSTM, and AE-LSTM) were validated on two case studies: a multi-degree-of-freedom BoucWen system and a 37-story fiber-discretized nonlinear steel moment-resisting frame, both subjected to stochastic seismic excitation and structural parameter uncertainty. All three approaches achieved low errors: MLPLSTM was most accurate for the lower-dimensional BoucWen system, whereas MPNNLSTM and AELSTM performed best on the more complex steel frame. In addition, predictive variance consistently tracked actual error, supporting active learning and confidence assessment in stochastic response trajectory prediction.

Haimiti Atila, Seymour Spence
University of Michigan
hatila@umich.edu, smjs@umich.edu

MS87

No-Replication Stochastic Emulation for Multi-Output Responses and Application to Seismic Risk Assessment

Seismic risk assessment is a promising application area of surrogate modeling, providing computationally efficient approximations of complex structural dynamic behavior. A challenge, however, is the presence of aleatoric uncertainty in surrogate predictions, which arises when embedding high-dimensional stochasticity of earthquake ground motions into the low-dimensional input of the surrogate. The resulting indeterminacy of structural response prompts the application of stochastic emulation, which can predict the entire probability distributions of model responses. However, challenges exist in its application: (i) often the training points are selected from existing ground motion databases, limiting the acquisition of replicated samples and restricting training data size; (ii) structural outputs tend to exhibit significant correlations, requiring a new methodology to faithfully capture the joint statistics of multi-output predictions. In response to these gaps, we propose a Gaussian process (GP)-based multi-output stochastic emulation method capable of predicting the joint probability distribution of outputs, while not relying on sample replications. This is achieved by introducing auxiliary surrogates that predict the second-order statistics of the outputs, leveraging observations of sample deviation statistics. The proposed framework can accommodate Gaussian or non-Gaussian marginal predictions obtained using existing stochastic emulation methods without extra simulations.

Sang-ri Yi
Rice University
sangri.yi@rice.edu

Alexandros A. Taflanidis
University of Notre Dame
a.taflanidis@nd.edu

MS87

Efficient Excursion Set Estimation for Noisy Models: A Revisit of Stepwise Uncertainty Reduction with New Insights

Excursion sets play a critical role in engineering as they typically represent the feasible parameters under constraints on the system performance. They appear naturally in reliability analysis, consisting of the parameters for which the system response exceeds a specified threshold. Over the past decades, active learning strategies for excursion set estimation of expensive-to-evaluate models have received considerable attention [Chevalier, Fast uncertainty reduction strategies relying on Gaussian process models]. In particular, stepwise uncertainty reduction (SUR) provides powerful sequential design strategies, tackling excur-

sion set estimation problems by aiming at the reduction of a global uncertainty measure [Bect et al., Sequential design of computer experiments for the estimation of a probability of failure]. When dealing with noisy or stochastic models, active learning remains challenging, as uncertainty in the observations complicates the estimation process. In this presentation, we revisit the general SUR framework for the estimation of excursion sets and its adaptation to noisy models. To improve the efficiency of SUR-based approaches, we incorporate recent developments in set inversion [Ait Abdelmalek-Lomenech et al., Bayesian Active Learning of (small) Quantile Sets through Expected Estimator Modification]. The performance and effectiveness of the proposed methods are illustrated through several numerical examples.

Romain Ait Abdelmalek-Lomenech
Université Paris-Saclay, CNRS, CentraleSupélec,
Laboratoire des Signaux et Systèmes
romain.ait@centralesupelec.fr

Julien Bect
Université Paris-Saclay, CNRS, CentraleSupélec
Laboratoire des signaux et systèmes
julien.bect@centralesupelec.fr

Emmanuel Vazquez
CentraleSupelec
emmanuel.vazquez@centralesupelec.fr

Xujia Zhu
Université Paris Saclay
CentraleSupelec
xujia.zhu@centralesupelec.fr

MS88

Localised Surrogate Transition: Fast-mixing MCMC Without Gradients

In this talk, I will revisit Multilevel Delayed Acceptance (MLDA) an efficient Markov chain Monte Carlo (MCMC) method for high-dimensional Bayesian inverse problems that does not require gradient information of the underlying posterior distribution. Instead, it uses a hierarchy of surrogates to accelerate convergence, e.g., coarser approximations of the likelihood. In previous works, we have demonstrated the fast convergence and the high effective sample size for MLDA, as well as multilevel variance reduction. In recent theoretical works, it has been shown that including gradient information allows to improve the upper bound on the total-variation mixing-time of MCMC from for Random Walk Metropolis to for MALA or for HMC. Here, is the condition number of a -strongly convex and -smooth posterior distribution. We prove that for a suitable hierarchy of surrogates the same speedup as for MALA can be achieved by MLDA without explicit gradient evaluations. This requires two modifications of the surrogate densities in MLDA, akin to Tikhonov regularisation and tempering. Beyond these theoretical results, our numerical experiments indicate that successive states of the chain exhibit substantially reduced autocorrelation, even compared to the original MLDA. Numerical evidence suggests that incorporating tempering into the surrogate densities also improves robustness with respect to multimodality of the target density.

Robert Scheichl
Ruprecht-Karls University Heidelberg
r.scheichl@uni-heidelberg.de

Robert Kutri
University of Heidelberg
robert.kutri@uni-heidelberg.de

MS89

Harnessing Huge Ensembles for Data-Driven Extreme Weather Prediction

FourCastNet 3 (FCN3) advances global weather modeling by implementing a scalable, geometric machine learning (ML) approach to probabilistic ensemble forecasting. The approach is designed to respect spherical geometry and to accurately model the spatially correlated probabilistic nature of the problem, resulting in stable spectra and realistic dynamics across multiple scales. FCN3 delivers forecasting accuracy that surpasses leading conventional ensemble models and rivals the best diffusion-based methods, while producing forecasts 8 to 60 times faster than these approaches. In contrast to other ML approaches, FCN3 demonstrates excellent probabilistic calibration and retains realistic spectra, even at extended lead times of up to 60 days. All of these advances are realized using a purely convolutional neural network architecture tailored for spherical geometry. Scalable and efficient large-scale training on 1024 GPUs and more is enabled by a novel training paradigm for combined model- and data-parallelism, inspired by domain decomposition methods in classical numerical models. Additionally, FCN3 enables rapid inference on a single GPU, producing a 60-day global forecast at 0.25, 6-hourly resolution in under 4 minutes. Its computational efficiency, medium-range probabilistic skill, spectral fidelity, and rollout stability at subseasonal timescales make it a strong candidate for improving meteorological forecasting and early warning systems through large ensemble predictions.

Boris Bonev
Nvidia
bbonev@nvidia.com

Thorsten Kurth
NVIDIA
tkurth@nvidia.com

Ankur Mahesh
Lawrence Berkeley National Laboratory
University of California Berkeley
amahesh@lbl.gov

Mauro Bisson, Marius Koch, Georg Ertl, Dallas Foster,
Jean Kossaifi, Karthik Kashinath
NVIDIA
mbisson@nvidia.com, mkoch@nvidia.com,
gertl@nvidia.com, dfoster@nvidia.com,
jkossaifi@nvidia.com, kkashinath@nvidia.com

Anima Anandkumar
California Institute of Technology & NVIDIA
anima@caltech.edu

William D. Collins
Lawrence Berkeley National Laboratory
wdcollins@lbl.gov

Mike Pritchard, Alex Keller
NVIDIA

mpritchard@nvidia.com, akeller@nvidia.com

MS89

Amortized Sequential Bayesian Inference and Experimental Design with Attention-Based Conditional Normalizing Flow

Sequential Bayesian inference involves recursively updating the probability distribution of the uncertain parameter in a model based on a stream of noisy and sparse observations of the environment. It plays a central role in problems cast as partially-observed Markov decision processes, e.g., digital twins and sequential Bayesian experimental design. In this talk, I introduce a method that uses a single attention-based conditional normalizing flow for the amortized approximation of the sequence of posteriors conditioned on streaming actiondata pairs (the observation history). This method has three key features. First, it can learn the evolution of parametric uncertainty by predicting a posterior given any observation history of arbitrary length, extrapolating well beyond the training horizon. Second, it is naturally permutation-invariant to the order of the observation history, while allowing the conditional structure of the likelihood to be encoded via positional embeddings. Third, it enables fast, stable, and explicit density evaluations once trained, suitable for downstream tasks such as estimating incremental expected information gains for sequential experimental design. I show promising numerical results on sequential inference problems such as signal-source localization finding, inference of patient-specific pharmacokinetics/dynamics, and inference of nonlinear viscoelasticity via controlled uniaxial tension.

Lianghao Cao
California Institute of Technology
lianghao@caltech.edu

MS89

Combining Inverse Methods and Machine Learning Through Differentiable Earth System Models

Despite the revolution in autonomous in-situ and satellite remote sensing, much of the ocean remains poorly sampled observationally in terms of its spatio-temporal evolution. Data assimilation seeks to combine the incomplete knowledge reservoirs from sparse observations and numerical models that require the provision of uncertain inputs to produce optimal state and/or parameter estimates of the system considered. In its various renderings, methods of optimal estimation and control are being used to learn uncertain parameters in models (calibration problem) or to produce optimal initial states for prediction (extrapolation problem). Scientific machine learning methods that learn neural networks or operators rely on similar mathematical concepts that underpin optimal estimation and may be successful when the system is heavily sampled. We explore the concept of differentiable programming in the context of ocean modeling. It is implemented in the context of general-purpose automatic differentiation applied to a new ocean model written in the Julia programming language that is augmented with ML. We provide a perspective on opportunities and challenges with a view toward hybrid data assimilation and machine learning to enable online or full-model learning.

Patrick Heimbach
Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin, Austin, TX
heimbach@oden.utexas.edu

William S. Moses
University of Illinois at Urbana-Champaign, U.S.
wsmoses@illinois.edu

Joseph L. Kump
University of Texas at Austin
josek97@utexas.edu

Sarah M. Williamson
University of Texas
swilliamson@utexas.edu

Dhruv Apte
The University of Texas at Austin
dhruv.apte@utexas.edu

Paul Berg
Bern University
paul@ber.gp

Valentin Churavy
Massachusetts Institute of Technology
vchuravy@mit.edu

Mosè Giordano
University College London
m.giordano@ucl.ac.uk

Chris Hill
Earth, Atmospheric and Physical Sciences
M.I.T.
cnh@mit.edu

Nora Loose

C

Worthy
nora.loose@gmail.com

Alexis Montoison, Shk Narayanan
Argonne National Laboratory
alexis.montoison@polymtl.ca, snarayan@mcs.anl.gov

Avik Pal
Massachusetts Institute of Technology
avikpal@mit.edu

Michel Schanen
Argonne National Laboratory
mschanen@anl.gov

Greg Wagner
Massachusetts Institute of Technology
wagner.greg@gmail.com

MS89

Can Transformers Discover Physical Laws from Data? Scaling General-Purpose Models from Molecular Dynamics to 3D Turbulence

Scientific machine learning often relies on rigid "physics-informed" architectures. However, as the broader deep learning community converges toward domain-agnostic Transformers, a fundamental question arises: must physical inductive biases always be engineered, or can they

emerge by learning from large-scale data? We investigate this across two regimes: atomistic modeling and 3D turbulence. First, we show unmodified Transformers achieve competitive accuracy in molecular prediction tasks without explicit physical priors. Analysis reveals that the models autonomously learn internal representations that align with physical structures like electrostatics. Second, we introduce EddyFormer, integrating spectral-element tokenization with a Transformer backbone for fluid dynamics. EddyFormer achieves a 30x acceleration over Direct Numerical Simulation at 256^3 resolution. Applied to unseen domains 4x larger than in training, EddyFormer demonstrates domain generalization, preserving physics-invariant metrics such as energy spectra, correlation and structure functions. Together, these results suggest that general-purpose architectures can autonomously discover complex physical structure from data, pointing toward flexible and scalable approaches for multi-scale modeling. We conclude by discussing how this framework can extend to low-data regimes, serving as flexible data-driven priors where physical constraints are strategically applied.

Aditi Krishnapriyan
Chemical Engineering and EECS, UC Berkeley
aditik1@berkeley.edu

MS90

Assimilative Causal Inference

Causal inference is fundamental across scientific disciplines, yet existing methods struggle to capture instantaneous, time-evolving causal relationships in complex, high-dimensional systems. In this paper, assimilative causal inference (ACI) is developed, which leverages Bayesian data assimilation to trace causes backward from observed effects. ACI solves the inverse problem rather than quantifying forward influence. It uniquely identifies dynamic causal interactions without requiring observations of candidate causes, accommodates short datasets, and scales efficiently to high dimensions. Crucially, it provides online tracking of causal roles, which may reverse intermittently, and facilitates a mathematically rigorous criterion for the causal influence range, revealing how far effects propagate. ACI opens new avenues for studying complex systems, where transient causal structures are critical.

Nan Chen, Nan Chen
University of Wisconsin - Madison
chennan@math.wisc.edu, chennan@math.wisc.edu

Marios Andreou
University of Wisconsin-Madison
mandreou@wisc.edu

Erik Bollt
Clarkson University
bolltem@clarkson.edu

MS90

Data-Driven Dynamical Precursors of Extreme Events in High-Dimensional Chaotic Dynamics

Extreme events in high-dimensional chaotic dynamical systems often arise from rare, transient instabilities that are difficult to anticipate from limited observations. As a result, purely statistical predictors trained on historical time series frequently fail to generalize in the tail regimes that dominate risk. We present a fully data-driven framework for long-lead prediction of intermittent extremes that constructs interpretable dynamical precursors by tracking internal transient instabilities of the system. Instability is quantified using finite-time Lyapunov exponentlike indicators evaluated in a low-dimensional, adaptively evolving subspace spanned by Optimal Time-Dependent modes, avoiding the prohibitive cost of full tangent-space computations. These instability-informed precursors are integrated into a Transformer-based forecaster to produce probabilistic long-horizon predictions of energy dissipation and threshold exceedances. Using two-dimensional Kolmogorov flow as a canonical model of intermittent turbulence, we demonstrate that dynamical precursors substantially improve extreme-event predictability and extend practical lead times relative to observable-based indicators, providing a scalable and uncertainty-aware approach for forecasting rare events in complex systems.

Eirini Katsidoniotaki
Massachusetts Institute of Technology
Massachusetts Institute of Technology
eirka289@mit.edu

MS90

Large Deviation Theory-Informed Importance Sampling for Rare Event Estimation and Control

Rare and extreme events like hurricanes, energy grid blackouts, dam breaks, earthquakes, and pandemics are infrequent but have severe consequences. Because estimating the probability of such events can inform strategies that mitigate their effects, scientists must develop methods to study the distribution tail of these occurrences. However, calculating small probabilities is hard, particularly when involving complex dynamics and high-dimensional random variables. In this talk, I will discuss our proposed method for the accurate estimation of rare event or failure probabilities for expensive-to-evaluate numerical models in high dimensions, and its application to rare event control. The proposed approach combines ideas from large deviation theory and adaptive importance sampling. The importance sampler uses a cross-entropy method to find an optimal Gaussian biasing distribution, and reuses all samples made throughout the process for both, the target probability estimation and for updating the biasing distributions. Large deviation theory is used to find a good initial biasing distribution through the solution of an optimization problem. Additionally, it is used to identify a low-dimensional subspace that is most informative of the rare event probability. We compare the method with a state-of-the-art cross-entropy-based importance sampling scheme using examples including a tsunami problem.

Shanyin Tong
Columbia University
tong3@sas.upenn.edu

MS91

Fast Transforms for Spectral Modeling of Gaussian Random Fields

The specification of a covariance function is of paramount importance when employing Gaussian process models, but the requirement of positive definiteness severely limits those used in practice. Designing flexible stationary covariance functions is, however, straightforward in the spectral domain, where one needs only to supply a positive and symmetric spectral density. We introduce an adaptive integration framework for efficiently and accurately evaluating covariance functions and their derivatives at irregular loca-

tions directly from any continuous, integrable spectral density. We show that this framework can be made highly scalable by leveraging the nonuniform fast Fourier transform (NUFFT) in one dimension, and a novel nonuniform fast Hankel transform (NUFHT) in higher dimensions. As a result, we demonstrate several orders of magnitude speedup compared to naive uniform quadrature approaches, allowing us to evaluate covariance functions from slowly decaying, singular spectral densities at millions of locations to a user-specified tolerance in seconds on a laptop.

Paul Beckman
University of Texas at Austin
paul.beckman@austin.utexas.edu

MS91

Uncertainty-Aware Ptychographic Reconstruction: A Robust, Cost-Aware Inference Framework

We build on normalizing-flow surrogates of the ptychographic posterior to deliver calibrated uncertainty in both magnitude and phase while capturing multi-modal solutions. Using an amplitude-based likelihood with TV regularization, flows match the reconstruction quality of advanced iterative solvers and expose where uncertainty concentrates (e.g., reduced scan overlap and phasemagnitude crosstalk). We further develop a robust version at substantially lower cost, making flow-based UQ practical for large ptychography workloads.

Zichao Di
Argonne National Lab
wendydi@mcs.anl.gov

MS91

Exact Gaussian Processes for Massive Datasets Via Non-Stationary Sparsity-Discovering Kernels

A Gaussian Process (GP) is a prominent mathematical framework for stochastic function approximation in science and engineering applications. Its success is primarily attributed to the GP's analytical tractability, robustness, and natural inclusion of uncertainty quantification. Unfortunately, the use of exact GPs is prohibitively expensive for large datasets due to their unfavorable numerical complexity of $O(N^3)$ in computation and $O(N^2)$ in storage, where N is the size of the dataset. All existing methods addressing this issue utilize some form of approximation, which can lead to inaccuracies in function approximations and often limit the user's flexibility in designing expressive kernels. Instead of inducing sparsity via datapoint geometry and structure, we propose to take advantage of naturally-occurring sparsity by allowing the kernel to discover—instead of induce—sparse structure. The premise of this work is that the data sets and physical processes modeled by GPs often exhibit natural or implicit sparsities, but commonly used kernels do not allow us to exploit such sparsity. The core concept of exact, and at the same time sparse GPs, relies on kernel definitions that provide enough flexibility to learn and encode not only non-zero but also zero covariances. This principle of ultra-flexible, compactly supported, and non-stationary kernels, combined with HPC, lets us scale exact GPs well beyond 5 million data points.

Marcus M. Noack
Lawrence Berkeley National Laboratory
MarcusNoack@lbl.gov

MS91

Physics-Informed Gaussian Process with Applications in ODE/PDE Parameter Estimation

Parameter estimation for nonlinear dynamic system models, represented by ordinary differential equations (ODEs) or partial differential equations (PDEs), using noisy and sparse experimental data is a vital task in many fields. We propose a fast and accurate method, physics-informed Gaussian process, for this task. Our method uses a Gaussian process model over system components, explicitly conditioned on the physics information that gradients of the Gaussian process must satisfy the ODE/PDE system. By doing so, we completely bypass the need for traditional numerical integration and achieve substantial savings in computational time. Our method is also suitable for inference with unobserved system components and provides uncertainty quantification. Our method is distinct from existing approaches as we provide a principled statistical construction under a Bayesian framework, which rigorously incorporates the ODE/PDE system through conditioning.

Shihao Yang
Georgia Tech
shihao.yang@isye.gatech.edu

MS92

An Occam's Razor Approach to Uncovering and Encoding Multivariable Structure in Kernel Learning

Emulating the input-output behavior of black-box models demands approximation methods that are both interpretable and expressive. Existing approaches involve tradeoffs: kernels and orthogonal bases offer interpretability but limited expressivity or scalability, while deep neural networks achieve expressiveness at the cost of interpretability. We aim to bridge this gap through total sensitivity kernels, a parameterized kernel family where each parameter is linked to the total sensitivity of a corresponding input. Automatic tuning of these parameters adapts the kernel to the multivariable structure of the target function, discovering which inputs matter and how they interact. Numerical results demonstrate that the method drives more efficient and accurate approximations while revealing interpretable structure.

John Darges
North Carolina State University
Department of Mathematics
jdarges@emory.edu

MS92

Training Data Distribution Optimization for Scientific Machine Learning

In scientific machine learning, models are often deployed with parameter values or boundary conditions that differ from those used in training. This talk presents the learning where to learn problem, which asks how to design a training data distribution that minimizes average prediction error across a family of deployment regimes. I will outline a theoretical analysis that explains how the training distribution influences deployment accuracy and motivates two adaptive algorithms based on bilevel optimization and alternating optimization in the space of probability mea-

sures. I will also describe discretized implementations that use either parametric distribution classes or nonparametric particle based gradient flows. Numerical results on learning functions and solution operators of partial differential equations show that the optimized training distributions outperform nonadaptive designs and improve robustness to distribution shift.

Nicolas Guerra
Cornell University
nmg64@cornell.edu

MS92

Autoencoders for Inverse Problems

In this work, we describe a new data-driven approach for inverse problems that exploits technologies from machine learning, in particular autoencoder network structures. We consider a paired autoencoder framework, where two autoencoders are used to efficiently represent the input and target spaces separately and optimal mappings are learned between latent spaces, thus enabling forward and inverse surrogate mappings. We focus on interpretations using Bayes risk and empirical Bayes risk minimization, and we provide various theoretical results and connections to existing works on low-rank matrix approximations. Similar to end-to-end approaches, our paired approach creates a surrogate model for forward propagation and regularized inversion. However, our approach outperforms existing approaches in scenarios where training data for unsupervised learning are readily available but training pairs for supervised learning are scarce. Furthermore, we show that cheaply computable evaluation metrics are available through this framework and can be used to predict whether the solution for a new sample should be predicted well.

Emma Hart
Emory University
ehart5@emory.edu

MS92

A Latent-Consistent Variational Autoencoder for Diversifying Generation

In this talk, we present a structured variational autoencoder (VAE) architecture designed to preserve latent consistency, namely, the latent-consistent variational autoencoder (LCVAE). Latent consistency is considered a core property of generative models for better generation diversity. While existing works have attempted to enforce it through soft constraints (i.e., introducing a penalty term in loss functions), the proposed model, LCVAE, exactly satisfies latent consistency through hard constraints by construction. LCVAE generalizes a popular generative model, Glow, and utilizes the matrix pseudoinverse. Additional reparameterization techniques stabilize the training process. The effectiveness of LCVAE is demonstrated on the MNIST dataset with the task of generating images of the number 5 without including them in training.

Auroni Hashim
NC State
ahashim3@ncsu.edu

Yeonjong Shin, Yeonjong Shin
North Carolina State University
yeonjong_shin@ncsu.edu, yeonjong_shin@ncsu.edu

Ichiro Takeuchi, Shotaro Tsutsui, Yanheng Gu
Nagoya University
ichiro.takeuchi@mae.nagoya-u.ac.jp,
tsutsui.shotaro.nagoyaml@gmail.com,
gu.yanheng.nagoyaml@gmail.com

MS93

Recent Progress in Sampling Via Diffusion Model, Fine Tuning and Stochastic Optimal Control

The algorithmic task of sampling from unnormalized density (or probability mass function if it is a discrete setup) has been studied for decades due to its importance. Diffusion model on the other hand is a much more recent technology that powers generative AI. Even more recently, stochastic optimal control started to provide an effective way for fine-tuning a generative model such as diffusion model. Can these contemporary technologies help sampling? This talk will briefly report several recent advances in this regard, including both learning-based and non-learning-based methods, and illustrate a pleasant possibility of alleviating metastability when sampling from multimodal distribution.

Molei Tao
Georgia Tech
mtao@gatech.edu

MS93

Free Energy Barriers in Bayesian Inverse Problems

Bayesian methods as parameter estimation provide us with a posterior distribution for the parameter of interest. Given a prior distribution and observed data, we can form a posterior. Often the resulting posterior distribution in not known in analytical form, and thus we rely on MCMC methods to compute point estimators as well as posterior spread estimation for uncertainty quantification. In this talk, we will discuss when local gradient based MCMC methods can fail to sample the target region. This is due to the occurrence of a free energy barrier, which we will introduce and show numerically how one can observe such a phenomenon. This is joint work with Prof. Sven Wang.

Maia Tienstra
University of Potsdam
tienstra@uni-potsdam.de

MS93

Mixing Time of the Proximal Sampler Via Strong Data Processing Inequality

Sampling is a fundamental algorithmic task in many applications. Sampling has intricate connections to optimization, for example via the perspective of sampling as optimization in the space of probability distributions. Many sampling algorithms are based on discretizing the Langevin dynamics, and they have mixing time guarantees which parallel the convergence guarantees from optimization. In this talk, we study the Proximal Sampler algorithm, which is an approximate proximal discretization of the Langevin dynamics. We show that the Proximal Sampler algorithm has mixing time guarantees that match the Langevin dynamics in continuous time, and parallel the convergence guarantees of the proximal gradient method from optimization. We highlight our proof technique which proceeds by establishing a strong data processing inequality for the algorithm in various statistical divergences.

Andre Wibisono

Yale University
andre.wibisono@yale.edu

MS93

High-Accuracy Sampling from Constrained Spaces with the Metropolis-Adjusted Preconditioned Langevin Algorithm

In this work, we propose a first-order sampling method called the Metropolis-adjusted Preconditioned Langevin Algorithm for approximate sampling from a target distribution whose support is a proper convex subset of $\mathbb{R}^d$. Our proposed method is the result of applying a Metropolis-Hastings filter to the Markov chain formed by a single step of the preconditioned Langevin algorithm with a metric $\mathcal{G}$, and is motivated by the natural gradient descent algorithm for optimisation. We derive non-asymptotic upper bounds for the mixing time of this method for sampling from target distributions whose potentials are bounded relative to $\mathcal{G}$, and for exponential distributions restricted to the support. Our analysis suggests that if $\mathcal{G}$ satisfies stronger notions of self-concordance introduced in Kook and Vempala (2024), then these mixing time upper bounds have a strictly better dependence on the dimension than when is merely self-concordant. We also provide numerical experiments that demonstrates the practicality of our proposed method. Our method is a high-accuracy sampler due to the polylogarithmic dependence on the error tolerance in our mixing time upper bounds.

Ashia Wilson
Massachusetts Institute of Technology
ashia07@mit.edu

MS94

Derivative-Informed Recurrent Neural Operator for Infinite-Dimensional Sequential Bayesian Optimal Experimental Design

We present a computational framework for sequential Bayesian optimal experimental design (SBOED) for PDE-constrained systems with infinite-dimensional uncertainties. We introduce an adaptive terminal optimality criterion that enables adaptive global optimality, and an equivalent, computation-friendly reformulation using Laplace and low-rank posterior approximations. To accelerate SBOED, we develop the derivative-informed latent attention neural operator (LANO), which combines (i) derivative-guided dimensionality reduction, (ii) attention to capture latent dynamics, and (iii) efficient latent-space training with projected Jacobians. LANO provides fast, accurate surrogates for parameter-to-observable maps and their Jacobians, and supports reduced-space computation of MAP points, eigenpairs, and posterior sampling. In a tumor-growth application optimizing MRI acquisition times, LANO outperforms two alternative neural surrogates and matches FEM accuracy for MAP and eigenvalues, while delivering an amortized 180 speed-up.

Phillip Si, Peng Chen, Jinwoo Go
Georgia Institute of Technology
psi6@gatech.edu, pchen402@gatech.edu, jgo@bnl.gov

MS94

A Control-Oriented Approach to Optimal Sensor Placement

We propose a control-oriented optimal experimental design (cOED) approach for linear PDE-constrained Bayesian inverse problems. In particular, we consider optimal control problems with uncertain parameters that need to be estimated by solving an inverse problem, which in turn requires measurement data. We consider the case where data is collected at a set of sensors. While classical Bayesian OED techniques provide experimental designs (sensor placements) that minimize the posterior uncertainty in the inversion parameter, these designs are not tailored to the demands of the optimal control problem. In the present control-oriented setting, we prioritize the designs that minimize the uncertainty in the state variable being controlled or the control objective. We propose a mathematical framework for uncertainty quantification and cOED for parameterized PDE-constrained optimal control problems with linear dependence to the control variable and the inversion parameter. We also present scalable computational methods for computing control-oriented sensor placements and for quantifying the uncertainty in the control objective. Additionally, we present illustrative numerical results in the context of a model problem motivated by heat transfer applications.

Madhusudan Madhavan
North Carolina State University
mmadhav@ncsu.edu

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Bart G. Van Bloemen Waanders, Rebekah White
Sandia National Laboratories
bartv@sandia.gov, rebwhit@sandia.gov

MS94

Goal Oriented Optimal Design of Infinite-dimensional Bayesian Inverse Problems Using Quadratic Approximations

We consider goal-oriented optimal experimental design (OED) for Bayesian linear inverse problems governed by partial differential equations (PDEs) with infinite-dimensional inversion parameters. In particular, we seek sensor placement strategies that minimize the posterior variance of a goal quantity of interest. The target quantity is assumed to be a nonlinear functional of the inversion parameters. We propose an OED approach that uses a quadratic approximation of the goal-functional to define our experimental design criterion which is obtained by integrating the quadratic approximation over the data. Under the assumption of Gaussian prior and noise models, we derive a closed-form expression for this criterion. Subsequently, we propose efficient and accurate computational methods and use a greedy approach to obtain our sensor placement strategy. We illustrate the proposed approach for two model inverse problems governed by PDEs. Our numerical results demonstrate the proposed approach outperforms non-goal-oriented (A-optimal) and linearization-based (c-optimal) approaches.

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

John N. Neuberger
North Carolina State University
jnneuber@ncsu.edu

MS95

Optimal Sensor Placement for Gaussian Processes Using Column Subset Selection

Gaussian processes use data measured at a finite number of sensor locations to reconstruct a spatially dependent function with quantified uncertainty. However, since one can only deploy a finite number of sensors, it is important to determine how to optimally place the sensors to minimize a measure of the uncertainty in the reconstruction. We consider the Bayesian D-optimality criterion to determine the optimal sensor locations by choosing k sensors from a candidate set of m sensors. Since this is an NP-hard problem, our approach uses sensor placement as a column subset selection problem (CSSP) on the covariance matrix, computed using the kernel function on the candidate sensor points. We propose an algorithm that uses the Golub-Klema-Stewart framework (GKS) to select sensors and provide an analysis of lower bounds on the D-optimality of these sensor placements. To reduce the computational cost in the GKS step, we propose and analyze randomized Nystrm approximations on the kernel matrix. We demonstrate the performance of our method on two applications: fiber coating dynamics and sea surface temperature.

Jessie Chen, Hangjie Ji, Arvind Saibaba
North Carolina State University
jchen228@ncsu.edu, hji5@ncsu.edu, asaibab@ncsu.edu

MS95

Amortized Inference and Dimensionality Reduction for Bayesian Inverse Problems on Function Spaces

We present algorithms and analysis for Bayesian inverse problems on function spaces via prior-preserving truncations. Given a posterior distribution $d\nu = \exp(-\Phi(u))d\mu$ on a Hilbert space $\mathcal{H}$, we consider approximate distributions of the form $d\hat{\nu} = \exp(-\phi \circ \psi(u))d\mu$, where $\psi : \mathcal{H} \mapsto \mathbb{R}^d$. We study its approximation properties of this class of measures and present efficient algorithms for computation.

Alexander Hsu, Bamdad Hosseini
University of Washington
owlx@uw.edu, bamdadh@uw.edu

Ricardo Baptista
University of Toronto
r.baptista@utoronto.ca

MS95

Bayesian Calibration and Optimization for Pulsed Power

Traditionally, there are two pillars of science: theory and experimentation. These two inform one another and lead scientists to make educated guesses and decisions toward advancing science. More recently, the driving force behind scientific advancement has not just focused on how much information can be learned, but how quickly. Additionally, experimental data can be costly and difficult to obtain. With these motivations in mind, the field of experimental design aims to maximize the information gained from as few experimental data points as possible. Computation, as the third pillar of science, has been used to facilitate optimal experimental design. Sandias Z machine is the worlds most powerful and efficient laboratory radiation source. A principal uncertainty in Z experiments is how effectively current can be delivered. Power flow simulations are very intensive, making them infeasible to use in critical design and optimization studies. Developing a consistent picture of how losses develop and evolve would improve understanding of present-day experiments and better constrain circuit model representations, providing a basis for quantifying uncertainties in circuit models applied to Z and improve confidence in predictions of target performance. This presentation details the implementation of a Bayesian optimization study to maximize the information gained from Z experimental data and design. * SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525

Kathryn Maupin, Christopher Jennings, Brian Hutsel
Sandia National Laboratories
kmaupin@sandia.gov, cjennin@sandia.gov, bthutse@sandia.gov

MS95

Learning Paths for Dynamic Measure Transport: A Regularity Control Perspective

Sampling from a target probability distribution is a fundamental task in data assimilation and inverse problems, wherein the target distribution of interest is a Bayesian posterior. Many recent sampling approaches make use of dynamics, such as ordinary or stochastic differential equations, to transport samples from a tractable reference distribution (such as the prior) to the target distribution. This dynamic measure transport (DMT) approach induces a path of distributions, corresponding to the distributions of the intermediate states of the dynamics, between the reference and the target, and in some DMT approaches it is possible to specify this path explicitly. Understanding of how to choose paths well, however, is lacking, and some paths, including the geometric annealing path commonly employed in Bayesian applications, can produce undesirable teleportation behavior. In this talk we bring a control perspective to the problem, posing a flexible family of optimization problems for identifying paths of distributions for DMT and advocating for the use of objective terms which encourage smoothness of the corresponding velocity fields. We present a numerical algorithm for solving these problems based on recent Gaussian process methods for solution of partial differential equations and demonstrate the ability of our method to recover more efficient and smooth transport models compared to those which use the geometric annealing path.

Aimee Maurais
Massachusetts Institute of Technology
maurais@mit.edu

Bamdad Hosseini
University of Washington
bamdadh@uw.edu

Youssef M. Marzouk
Massachusetts Institute of Technology

ymarz@mit.edu

MS96

Scalable Given-Data Sobol Index Estimators for Very High-Dimensional Sensitivity Analysis in Neural Networks

Commonly used methods for Global Sensitivity Analysis (GSA), such as the pick-freeze Saltelli approach for computing Sobol' indices, require a number of samples that scales linearly with the number of parameter inputs, and a forward model that can evaluate samples at arbitrary locations. In this work, we make use of given-data methods for GSA and extend them to scale up to problems with extremely high-dimensional input spaces. Our extensions include a general definition of the given-data Sobol index estimator with arbitrary partition, a streaming algorithm to process input-output samples in batches, and a heuristic to filter out small indices that are indistinguishable from zero indices due to statistical sampling noise. We apply our approach to compute the sensitivity of neural network outputs with respect to the neural network weights, and demonstrate scaling up to hundreds of thousands of inputs.

Teresa Portone
Sandia National Laboratories
tporton@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Samantha Yang, Emiliano Islas-Quinones, Patrick Xiao
Sandia National Labs
ysamantha.mail@gmail.com, eislasq@arizona.edu, txiao@sandia.gov

MS96

Diffusion Non-Additive Model for Multi-Fidelity Simulations with Tunable Precision

Computer simulations are indispensable for analyzing complex systems, yet high-fidelity models often incur prohibitive computational costs. Multi-fidelity frameworks address this challenge by combining inexpensive low-fidelity simulations with costly high-fidelity simulations to improve both accuracy and efficiency. However, certain scientific problems demand even more accurate results than the highest-fidelity simulations available, particularly when a tuning parameter controlling simulation accuracy is available, but the exact solution corresponding to a zero-valued parameter remains out of reach. In this paper, we introduce the Diffusion Non-Additive (DNA) model, inspired by generative diffusion models, which captures nonlinear dependencies across fidelity levels using Gaussian process priors and extrapolates to the exact solution. The DNA model: (i) accommodates complex, non-additive relationships across fidelity levels; (ii) employs a nonseparable covariance kernel to model interactions between the tuning parameter and input variables, improving both predictive performance and physical interpretability; and (iii) provides closed-form expressions for the posterior predictive mean and variance, allowing efficient inference and uncertainty quantification. The methodology is validated on a suite of numerical studies and real-world case studies. An R package implementing the proposed methodology is available to support practical applications.

Junoh Heo, Romain Boutelet
Michigan State University
heojunoh@msu.edu, boutelet@msu.edu

Chih-Li Sung
Michigan State University
Department of Statistics and Probability
sungchih@msu.edu

MS96

Diffusion Models for Bayesian Inverse Problems in Low Temperature Plasmas

A common task across all fields of science and engineering is to predict the inputs of a model or simulation given partial observations of the models outputs. These "inverse problems" are notoriously difficult to solve. Classical optimization methods are commonly used but only provide one answer and thus cannot quantify the uncertainty in that answer. Bayesian inference methods, such as Markov chain Monte Carlo (MCMC), address this need by generating samples from the posterior distribution of the inputs, specified by a likelihood function. The generated samples can be analyzed to quantify uncertainty in model predictions. Such methods are robust and widely used, but are difficult to parallelize and have challenges scaling to high dimensions. To this end, we propose and demonstrate the use of a diffusion model for solving inverse problems in the field of low temperature plasmas. These models have become extremely widely-used for generative modeling tasks, particularly image generation. They have also been applied more recently to text generation and the solution of inverse problems in science and engineering. Specifically, we apply a diffusion model to the task of learning unknown scalar input fields given partial observation of the results of a 1-D fluid simulation of a Hall thruster. We show that these methods provide a powerful, robust, and fast alternative to traditional Bayesian inference methods.

Thomas Marks, Alex Gorodetsky
University of Michigan
marksta@umich.edu, goroda@umich.edu

MS96

Specifying Prior Distributions in Reliability Applications

In reliability analysis, data are often sparse due to limited failure events or complex censoring mechanisms, making traditional inference challenging. Bayesian methods offer a principled framework for incorporating prior knowledgesuch as physics-of-failure models or historical datainto the analysis, enabling more robust inference under data scarcity. This paper explores strategies for specifying prior distributions in reliability contexts, emphasizing both informative and default priors. It reviews existing approaches that yield desirable frequentist properties, including coverage accuracy, and introduces extensions tailored to the practical realities of reliability data. Through theoretical discussion and applied examples, the paper illustrates how carefully chosen priors can enhance inference quality, support decision-making, and address challenges posed by censoring and limited failure observations.

Qinglong Tian
University of Waterloo
qinglong.tian@uwaterloo.ca

Colin Lewis-Beck, Jarad Niemi, William Meeker
Iowa State University
clewisbe@iastate.edu, niemi@iastate.edu, wqmeeker@iastate.edu

MS97

A Preliminary Phase Space Decomposition of the Three-Body Model for Long-Time Uncertainty Propagation

Long-time uncertainty propagation methods for unperturbed and weakly perturbed two-body models have been developed for the purpose of predicting the debris environment at near-Earth orbits over decade long time spans. These methods often leverage averaging techniques, ergodicity theory, and a dissipative factor due to the presence of the Earth's upper atmosphere. Equivalent long duration models for a three-body environment, such as the cislunar domain, do not yet exist. In this talk we introduce a possible path forward to achieve analogous long-time uncertainty propagation in the three-body model using the fact that periodic orbits are dense to generate a finite covering and hence decomposition of the phase space. From this preliminary phase space decomposition we present initial investigations on uncertainty representation and transfer from localized neighborhoods of the model to globalization. Techniques from normal form theory and Markov processes are used in the definition of the decomposition and modeling external as well as truncation terms for uncertainty propagation.

Jannik Graebner, Jack Schwartz, Ryne Beeson
Princeton University
jg3607@princeton.edu, js1765@princeton.edu, ryne@princeton.edu

MS97

Uncertainty Propagation in Orbital Dynamics with Arbitrary Probability Distribution and Quantified Error Bounds

The FokkerPlanck partial differential equation (FP-PDE) governs the time evolution of uncertainty in stochastic dynamical systems. In orbital dynamics, solving the FP-PDE is challenging due to nonlinear motion, high-dimensional states, and large spacetime domain. We present a physics-informed neural network (PINN) approach that approximates the FP-PDE solution—the probability density over continuous spacetime $p(x, t)$. This approach is agnostic to coordinates and neural network designs. We outline relevant coordinate choices and introduce architectures that enforce probability-density structure, supporting efficient sampling or analytic marginalization. We analyze approximation error and derive upper bounds on average and worst-case deviations—total variation and L_∞, respectively—in terms of the training loss. These certificates support tasks such as bounding conjunction probabilities and performing statistical orbit determination from ensembles of probability densities. Numerical studies on two-body dynamics show accurate propagation of non-Gaussian uncertainty, tight error bounds, and practical training. Although evaluated on the two-body problem, the framework generalizes to more complex regimes, including cislunar space, where orbital uncertainty propagation is critical. More broadly, it applies to well-posed FP-PDE problems and can benefit from advances in neural architectures (e.g., transformers) while retaining error certificates.

Chun-Wei Kong
University of Colorado Boulder
chun-wei.kong@colorado.edu

MS97

Approximate Transformation of Gaussian Moments Using the Canonical Polyadic Decomposition

Statistical moment analysis provides an efficient method to propagate uncertainty in dynamical systems. Covariance propagation with linearized dynamics is employed extensively in the analysis of uncertainty in spacecraft trajectories; however, in more nonlinear systems, the linearized dynamics approximation can lead to unacceptable errors in the final covariance. Additionally, since nonlinear dynamics can lead to non-Gaussian distributions, higher than second-order statistical moments may be desirable to quantify the non-Gaussian aspects of the distribution. Higher-order moment propagation methods employing computational Taylor series based approaches such as differential algebra or state transition tensors can be limited by their increased computational costs. We present a technique for higher-order moment propagation of an arbitrary initial Gaussian distribution by employing the canonical polyadic decomposition of the moment tensors associated with a Gaussian distribution. With the Gaussian moment tensors written as the sum of rank-1 tensors, propagation of moments using higher-order Taylor series becomes substantially less computationally intensive, decreasing the polynomial order of the computation complexity by at least half. The canonical polyadic decompositions of the standard Gaussian moments are expensive to compute, but need only be performed offline once and then can be easily transformed to find the decomposition for an arbitrary Gaussian.

Jackson Kulik, Braden Hastings
Utah State University
jackson.kulik@usu.edu, braden.hastings@usu.edu

MS97

Nonlinearity and Uncertainty Informed Gaussian Mixture Splitting for Astrodynamics Uncertainty Propagation

As strong space domain awareness capabilities become increasingly important, accurate representations of uncertainty propagated through nonlinear astrodynamical systems are essential. Gaussian mixtures can characterize the resulting non-Gaussian uncertainty while retaining the computationally convenient properties of Gaussians. Adaptive Gaussian mixture approaches increase resolution by splitting to account for dynamical nonlinearities that can induce significant non-Gaussianity upon the uncertainty distribution. To obtain an accurate yet tractable representation of uncertainty, the splitting direction should be chosen to best approximate the true propagated density with a limited number of mixands. This talk discusses optimal Gaussian splitting using higher-order measures of nonlinearity that also account for the underlying distribution through input whitening, uncertainty scaling, and spherical averaging. These heuristics utilize higher-order tensor approximations of the nonlinear solution flow, which are additionally exploited opportunistically to obtain higher-order approximations of the propagated moments. A deferred splitting technique is presented that

reduces the required number of numerical integration operations by leveraging the cocycle conditions of state transition tensors. These splitting heuristics are evaluated over three orbits and shown to accurately capture the transformed uncertainty distributions, with deferred splitting significantly improving runtimes.

Gerard A. Siciliano, Keith LeGrand
Purdue University
gsicilia@purdue.edu, klegrand@purdue.edu

Jackson Kulik
Utah State University
jackson.kulik@usu.edu

MS98

Regularized Wasserstein Proximal Algorithms for Nonsmooth Sampling Problems

In this talk, we introduce regularized Wasserstein proximal algorithms for nonsmooth sampling problems. Our approach is based on a splitting scheme for the time-implicit discretization of the probability flow ODE. In this framework, the score functiondefined as the gradient of the log-density of the current distributionis approximated via the regularized Wasserstein proximal. We establish convergence to the target distribution under suitable conditions and demonstrate the effectiveness of the method through numerical experiments on high-dimensional nonsmooth sampling tasks. We further highlight its application to generative modeling, including the design of a specialized class of transformers.

Fuqun Han
University of California, Los Angeles
fqhan@math.ucla.edu

MS98

A Meta-Learning Perspective on Transformer-Based Language Models and Their Learning Dynamics

The Transformer architecture has become prominent in developing large causal language models. However, mechanisms to explain its capabilities are not well understood. Focused on the training process, here we establish a meta-learning view of the Transformer architecture when trained for the causal language modeling task, by explicating an inner optimization process within the Transformer. Further, within the inner optimization, we discover and theoretically analyze a special characteristic of the norms of learned token representations within Transformer-based causal language models. Our analysis is supported by experiments in various settings.

Xinbo Wu
University of Illinois at Urbana-Champaign
xinbowu2@illinois.edu

MS99

A Variational Approach for Sparse Identification of Non-Local Potential in Aggregation-Diffusion Equations

We address the inverse problem of identifying nonlocal interaction potentials in nonlinear aggregation-diffusion equations from noisy discrete trajectory data. Our approach involves formulating and solving a regularized variational problem, which requires minimizing a quadratic error functional across a set of hypothesis functions, further augmented by a sparsity-enhancing regularizer. We employ a partial inversion algorithm, akin to the CoSaMP and subspace pursuit algorithms, to solve the Basis Pursuit problem. A key theoretical contribution is our novel stability estimate for the PDEs, validating the error functional ability in controlling the 2-Wasserstein distance between solutions generated using the true and estimated interaction potentials. Our work also includes an error analysis of estimators caused by discretization and observational errors in practical implementations. We demonstrate the effectiveness of the methods through various 1D and 2D examples showcasing collective behaviors.

Sui Tang
University of California Santa Barbara
suitang@math.ucsb.edu

MS99

Data-Driven Enhanced Variational Modeling of Flocking Dynamics

Variational models play a central role in physics, materials science, biology, and data science. Such models are governed by a free energy functional and rate of energy dissipation. A key challenge in building variational models lies in determining these two quantities from data. In this talk, we present a PDE-free learning framework that separately identifies the free energy and the dissipation mobility directly from simulation data. The method is based on the minimizing movement scheme, which naturally encodes the energydissipation law in a discrete-time formulation. As a case study, we demonstrate the framework on several illustrative examples, including flocking dynamics.

Yiwei Wang
UC Riverside
yiweiw@ucr.edu

MS100

Mowe : A Mixture of Weather Experts

Data-driven weather models have recently achieved state-of-the-art performance, yet progress has plateaued in recent years. This paper introduces a Mixture of Experts (MoWE) approach as a novel paradigm to overcome these limitations, not by creating a new forecaster, but by optimally combining the outputs of existing models. The MoWE model is trained with significantly lower computational resources than the individual experts. Our model employs a Vision Transformer-based gating network that dynamically learns to weight the contributions of multiple "expert" models at each grid point, conditioned on forecast lead time. This approach creates a synthesized deterministic forecast that is more accurate than any individual component in terms of Root Mean Squared Error (RMSE). Our results demonstrate the effectiveness of this method, achieving up to a 10% lower RMSE than the best-performing AI weather model on a 2-day forecast horizon, significantly outperforming individual experts as well as a simple average across experts. This work presents a computationally efficient and scalable strategy to push the state of the art in data-driven weather prediction by making the most out of leading high-quality forecast models.

Dibyajyoti Chakraborty, Romit Maulik
Pennsylvania State University
d.chakraborty@psu.edu, rmm7011@psu.edu

Sanjay Choudhry, Peter Harrington, Mohammad A. Nabian
Nvidia
schoudhry@nvidia.com, pharrington@nvidia.com, mnabian@nvidia.com

Dallas Foster
NVIDIA
dfoster@nvidia.com

MS100

Direct Regret Optimization in Bayesian Optimization

Bayesian optimization is a powerful paradigm for optimizing expensive black-box functions. Traditional BO methods typically rely on separate hand-crafted acquisition functions and surrogate models for the underlying function and often operate in a myopic manner. We propose a novel direct regret optimization approach that jointly learns the optimal model and non-myopic acquisition by distilling from a set of candidate models and acquisitions, and explicitly targets minimizing the multi-step regret. Our framework leverages an ensemble of Gaussian Processes to generate simulated BO trajectories, each guided by an acquisition function chosen from a pool of conventional choices, until a Bayesian early stop criterion is met. These simulated trajectories, capturing multi-step exploration strategies, are used to train an end-to-end decision transformer that directly learns to select next query points aimed at improving the ultimate objective. Experimental results on synthetic and real-world benchmarks suggest that our method consistently outperforms BO baselines, achieving lower simple regret and demonstrating more robust exploration in high-dimensional or noisy settings.

Fengxue Zhang, Yuxin Chen
The University of Chicago
zhangfx@uchicago.edu, chenyuxin@uchicago.edu

MS100

Scalable and Practical Bayesian Deep Learning with Variational Last Layers: Methods and Applications

Bayesian deep learning methods have had minimal impact on the largest, most impactful neural network models. In this talk, I will describe progress on highly scalable Bayesian methods that can be easily built into existing architectures. By focusing only on uncertainty in the last (few) layers, combined with a minimal set of maximally impactful architectural interventions, we achieve scalable and practical uncertainty quantification methods that are well suited to modern neural network design. I will describe use of these models in active learning tasks including Bayesian optimization, showing state of the art performance on various scientific problems.

James Harrison
Google DeepMind
jamesharrison@google.com

MS100

Covering Multiple Objectives with a Small Set of Solutions Using Bayesian Optimization

In multi-objective black-box optimization, the goal is typically to find solutions that optimize a set of T black-box objective functions simultaneously. Traditional approaches often seek a single Pareto-optimal set that balances trade-offs among all objectives. In our recent work, we consider a problem setting that departs from this paradigm: finding a smaller set of K solutions, where K ¡ T, that collectively "covers" the T objectives. A set of solutions is defined as "covering" if, for each objective, there is at least one good solution. A motivating example for this problem setting occurs in drug design. For example, we may have T pathogens and aim to identify a set of K ¡ T antibiotics such that at least one antibiotic can be used to treat each pathogen. To address this problem, we propose Multi-Objective Coverage Bayesian Optimization (MOCOBO), a principled algorithm designed to efficiently find a covering set of solutions.

Natalie T. Maus
University of Pennsylvania
nmaus@seas.upenn.edu

MS101

Unsupervised In-context Operator Learning on the Probability Measure Space

Many fundamental problems on probability measure spaces, such as optimal transport, mean field games, and Wasserstein gradient flows, are computationally demanding, while existing learning-based methods often rely on single instance solvers and require costly retraining for each new instance. In context learning with transformer models offers a new paradigm for approximating families of operators from a few context examples without task specific retraining. I will discuss our recent work on an unsupervised in context operator learning framework for learning solution operators on probability measure spaces. The method is discretization free, making it effective for high dimensional measure transport problems, and does not require supervised solution labels, substantially reducing the cost of data generation. If time permits, I will also discuss a generalization error analysis of the proposed transformer-based model, connecting it to emerging theory on in context learning and highlighting broader theoretical implications.

Rongjie Lai
Purdue University
lairj@purdue.edu

MS102

Temporal Generative Modeling for Weather and Climate Prediction

Accurate and reliable weather and climate prediction is vital for societal planning, disaster mitigation, and understanding climate change. The increasing frequency of extreme weather events demands robust probabilistic forecasts that effectively quantify uncertainty. While generative models, particularly diffusion models, offer a promising path, methodological innovations are needed to overcome crucial limitations in their computational and data efficiency for these large-scale problems. In this talk, I will discuss methodological advancements that address these efficiency bottlenecks by explicitly leveraging the temporal structure of the data to design more efficient generative models for forecasting and emulation problems.

Salva Cachay
UC San Diego

salvaruehling@gmail.com

MS102

Quantifying Epistemic Uncertainty in Diffusion Models

Diffusion models generate diverse samples by design, but that diversity is not the same as model uncertainty. In practice, many uncertainty quantification (UQ) methods for diffusion models conflate aleatoric variability from the reverse diffusion noise with epistemic uncertainty arising from uncertain parameters. In this talk, we present a new methodology for estimating epistemic uncertainty in diffusion models (see https://arxiv.org/abs/2602.09170). Our approach is based on a FisherLaplace formulation that propagates parameter uncertainty through the reverse diffusion process. To make our approach scalable, we introduce FLARE (FisherLaplace Randomized Estimator), which approximates the Fisher information by uniformly subsampling parameters across all layers. Across synthetic time-series generation benchmarks designed to probe multimodality, mode switching, and low-density regions, FLARE produces more faithful epistemic uncertainty scores and enables more reliable filtering of low-confidence samples than BayesDiff and common Laplace-based baselines.

Adityi Gupta
University of California, Berkeley
agupta@icsi.berkeley.edu

MS102

Conditional Flow Matching for Data Assimilation in Chaotic and Turbulent Systems

Flow-based generative models trained via flow matching are emerging as expressive priors for scientific inverse problems. This talk studies training-free conditional generation, guiding a pretrained flow model at inference time to assimilate sparse, noisy observations without any retraining. We develop a simple taxonomy and benchmark of guidance schemes that (i) modify the source term of the learned flow with likelihood or physics gradients (observation and PDE-residual guidance) and (ii) apply iterative corrections during integration (predictorcorrector, score fusion, and hybrid schedules). On canonical chaotic and turbulent testbeds, we evaluate accuracy and efficiency using RMSE, calibration, spectral recovery, energy-budget drift, and wall-time. Our analysis provides a comprehensive comparison of their accuracy and computational efficiency, offering practical insights for applying conditional generative models to challenging scientific inverse problems. The practical recipes (guidance strength and scheduling vs. noise level, sensor layout, and solver discretization) distilled from these analyses will let practitioners turn any pretrained flow-matching model into a plug-and-play data assimilation engine for complex spatiotemporal physics.

Meet Parikh
Cornell University
mhp66@cornell.edu

Yaqin Chen
Notre Dame
ychen234@nd.edu

Jianxun Wang
Cornell University
jw2837@cornell.edu

MS103

Bi-Fidelity Diffusion Modeling for Rare-Event Probability Estimation

Estimating the probability of rare but critical system events is challenging when input dimensionality is high and high-fidelity simulations are costly. We present a bi-fidelity framework that directly models the low-dimensional quantity of interest (QoI) to compute probability of failure, without constructing a full input-to-output surrogate. Low-fidelity simulations provide a computationally inexpensive source of data for training a score-based diffusion model, which is augmented with sparse high-fidelity evaluations to accurately capture the QoI distribution. High-fidelity data are incorporated via an energy network that biases the diffusion model toward high-fidelity behavior. Energy-guided sampling efficiently explores rare-event regions, and tail-focused optimization ensures an accurate description of extreme events. This approach bypasses the curse of dimensionality inherent in traditional surrogates, efficiently capturing heavy tails and multi-modal features in high-dimensional systems, with applications in safety-critical engineering scenarios.

Audrey Gaymann
University of Colorado Boulder
audrey.gaymann@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS103

Quantifying Resource Allocation Uncertainty for Multi-Fidelity Estimation

This talk is about estimating a parameter θ from models of different fidelities. Each model has an evaluation cost. Our goal is to allocate a budget to model evaluations in order to construct an unbiased estimator of θ having minimal variance. Approximate control variates and the multilevel best linear unbiased estimator are frameworks that address this problem, but they assume the covariance between models is known. In practice, the covariance must be estimated empirically. We quantify and mitigate the effect of this estimation on the allocation of budget to model evaluations.

Aurya Javeed
Sandia National Laboratories
asjavee@sandia.gov

MS104

Surrogate-Based Probabilistic Fatigue Assessment of Transmission Lines with Gaussian Process Regression: From Damage Estimation to Inspection Planning

Transmission line towers are repeatedly subjected to fluctuating wind loads that drive fatigue damage in bolted connections, threatening long-term reliability. This study develops a surrogate-based probabilistic framework to assess fatigue damage and guide inspection planning. High-fidelity finite element simulations are performed to capture joint-level responses under both preloaded and loos-

ened bolt conditions, with fatigue quantified using rainflow cycle counting and Miners rule. The resulting dataset is used to train Gaussian Process Regression (GPR) surrogates, which provide rapid and accurate predictions of fatigue accumulation across diverse wind speeds and directions. By coupling these surrogates with joint probability distributions of typical winds and a Poisson-based extreme-event model, the framework synthesizes annual damage and propagates uncertainty through sampling-based simulations. Results highlight fatigue-critical connections, reveal the relative roles of typical versus extreme winds, and quantify the influence of stress concentration factors and bolt preloading states on vulnerability. Building on these insights, inspection strategies are optimized by incorporating detection probabilities, repair resets, and costrisk trade-offs. The proposed methodology delivers a computationally efficient, reliability-informed tool for fatigue risk assessment and inspection scheduling of transmission line systems.

Pooria Mazaheri
Ph.D. student
Iowa state university
mazaheri@iastate.edu

Alice Alipour
Iowa State University
alipour@iastate.edu

MS104

Dimensionality Reduction-Based Stochastic Simulator for Seismic Risk Analysis of Bridge Structures

Seismic risk analysis of bridge structures often requires high-dimensional stochastic simulations that couple structural models with ground motion uncertainties. However, the associated computational cost can be prohibitive, especially when quantifying low-probability responses or conducting parametric studies. This study introduces a dimensionality reduction-based stochastic simulator that efficiently captures the essential variability of high-dimensional input spaces while preserving the underlying stochastic structure of seismic excitations. The proposed framework combines dimensionality reduction techniques with a probabilistic sampling strategy to generate surrogate stochastic responses, thereby alleviating the need for repeated full-scale simulations. Application to a benchmark bridge model demonstrates that the proposed simulator can accurately reproduce key seismic risk metrics, including response distributions, tail characteristics, and sensitivity indices, with significantly reduced computational effort. These results highlight the potential of the proposed stochastic simulator as a practical tool for seismic risk assessment of complex bridge structures, enabling more efficient integration of structural reliability and resilience considerations in performance-based design.

Jungho Kim
Sejong University
junghokim@sejong.ac.kr

Ziqi Wang
University of California, Berkeley
ziqiwang@berkeley.edu

MS104

Introducing Bayesian Inference for Stochastic Polynomial Chaos Expansions

Stochastic simulators arise in engineering and the sciences whenever a model produces different outputs for the same inputs, due to intrinsic randomness. Their analysis is computationally demanding, as their output is a random variable requiring characterization of its conditional distribution. Surrogate models have thus been developed to emulate such simulators efficiently. Among them, stochastic polynomial chaos expansions (SPCE) stand out by relying on minimal distributional assumptions. SPCE introduces a latent variable to capture intrinsic randomness and builds a polynomial chaos expansion in the joint inputlatent space. However, its current training relies on multi-step sparse regression and cross-validation, which can be computationally expensive and prone to suboptimal models. We propose a Bayesian framework for SPCE that overcomes these limitations. Sparsity-inducing priors on the expansion coefficients enable principled basis selection, while priors on the noise variance provide regularization and remove the need for cross-validation. Posterior inference, performed via sampling methods such as Hamiltonian Monte Carlo, further yields predictive uncertainty estimates, a feature particularly valuable for active learning.

Arthur Hatstatt, Maliki Moustapha
ETH Zurich
Chair of Risk, Safety and Uncertainty Quantification
arthur.hatstatt@ibk.baug.ethz.ch,
moustapha@ibk.baug.ethz.ch

Stefano Marelli
Chair of Risk, Safety and Uncertainty Quantification
Institute of Structural Engineering, ETH Zurich
marelli@ibk.baug.ethz.ch

Bruno Sudret
ETH Zurich
Chair of Risk, Safety and Uncertainty Quantification
sudret@ibk.baug.ethz.ch

MS104

Uncertainty-Aware Fragility Modeling of Urban Building Exteriors Subjected to Hurricane-Induced Windborne Debris with Conditional Generative Adversarial Nets

Building exteriors are highly susceptible to windborne debris damage during hurricanes. Fragility functions that estimate the probability of exterior failure as a function of hurricane intensity are commonly used to assess this risk. However, developing the fragility function usually requires numerous expensive physics-based simulations. To address this challenge, this study proposes a new machine learning-based surrogate model that leverages conditional generative adversarial networks (cGAN) to model the distribution of the stochastic output, i.e., the maximum impact of windborne debris on the building exterior, conditioned on hurricane intensity. To ensure the accuracy and reliability of cGAN, a stochastic simulation-guided training strategy is proposed that treats the associated data preparation process as an augmented stochastic simulation problem. This approach enhances exploration of the input space and enables the formulation of a novel cGAN training metric tailored toward fragility modeling. The trained cGAN model can rapidly generate extensive damage realizations under various hurricane scenarios, which facilitates the construction of fragility functions. Further, Monte Carlo dropout is introduced to quantify the uncertainty of the fragility function established using cGAN. Such uncertainty pro-

vides end users with a measure of confidence in the derived fragility estimates. An illustrative example is presented to validate the efficacy of the proposed approach.

Ziluo Xiong, Gaofeng Jia
Colorado State University
ziluo.xiong@colostate.edu, gaofeng.jia@colostate.edu

Yue Dong
The Hong Kong Polytechnic University
yue24.dong@polyu.edu.hk

Yanlin Guo
Colorado State University
yanlin.guo@colostate.edu

MS105

Simulation and Data Assimilation in An Idealized Coupled Atmosphere-Ocean-Sea Ice Floe Model with Cloud Effects

Sea ice plays a crucial role in the climate system, particularly in the Marginal Ice Zone (MIZ), a transitional area consisting of fragmented ice between the open ocean and consolidated pack ice. As the MIZ expands, understanding its dynamics becomes essential for predicting climate change impacts. However, the role of clouds in these processes has been largely overlooked. This paper addresses that gap by developing an idealized coupled atmosphere-ocean-ice model incorporating cloud and precipitation effects, tackling both forward (simulation) and inverse (data assimilation) problems. Sea ice dynamics are modeled using the discrete element method, which simulates floes driven by atmospheric and oceanic forces. The ocean is represented by a two-layer quasi-geostrophic (QG) model, capturing mesoscale eddies and ice-ocean drag. The atmosphere is modeled using a two-layer saturated precipitating QG system, accounting for variable evaporation over sea surfaces and ice. Cloud cover affects radiation, influencing ice melting. The idealized coupled modeling framework allows us to study the interactions between atmosphere, ocean, and sea ice floes. Specifically, it focuses on how clouds and precipitation affect energy balance, melting, and freezing processes. It also serves as a testbed for data assimilation, which allows the recovery of unobserved floe trajectories and ocean fields in cloud-induced uncertainties.

Changhong Mou
Utah State University
changhong.mou@usu.edu

MS105

Data-Driven Representations of Model-Form Error in Oscillatory Systems

Many physical systems can be modeled as forced oscillators, such as the vertical motion of bridges or driving cars, or the strings of a musical instrument. In particular, we consider a critical oscillation in the field of space weather, i.e., fluctuations of the total electron content of the ionosphere under solar forcing. However, this system poses many modeling challenges, including (i) large observational errors; (ii) significant model form error (MFE) in reduced models; and (iii) the sheer scale of the computation. Therefore, we propose a data assimilation method that estimates the state and the MFE simultaneously. Given that there is large error in both observations and the model, we must take advantage of some known or assumed error structure in either or both of these sources. We demonstrate the method with the Duffing oscillator and show some preliminary results using real data from the upper atmosphere.

Noah Peterson
University of Colorado Boulder
noah.peterson-2@colorado.edu

MS105

A Labeled Random Finite Set Approach to Dynamical Model Identification

Traditional model and system identification techniques do not distinguish between uncertainty in parameters and uncertainty in model form. We aim to make use of the labeled random finite set framework (traditionally used for quantifying uncertainty about space objects) in order to decouple these two notions from each other. This allows us to model uncertainty about terms that might or might not exist but would otherwise have large parameters. We pose this as a new type of Bayesian parameter-and-model-form Bayesian estimation problem and extend canonical ensemble Kalman filtering techniques to this new approach.

Andrey A. Popov
University of Hawai'i at Manoa, U.S.
apopov@hawaii.edu

Alexander Apo
University of Hawaii at Manoa
aapo9@hawaii.edu

Andrei A. Klishin
University of Washington
aklishin@hawaii.edu

Andrei A. Klishin
University of Hawai?i at Manoa
aklishin@hawaii.edu

MS105

Modeling Partially Observed Nonlinear Dynamical Systems and Efficient Data Assimilation Via Conditional Gaussian Koopman Network

A discrete-time conditional Gaussian Koopman network (CGKN) is developed to learn surrogate models for efficient state forecasting and data assimilation (DA) in high-dimensional, complex dynamical systems. Focusing on nonlinear, partially observed systems common in engineering and Earth science, the approach leverages Koopman embedding to construct latent variables representing unobserved states, whose dynamics are conditionally linear given the observed states. This structure yields a conditional Gaussian system, enabling the posterior distribution of latent states to be evaluated analytically. These closed-form DA updates are embedded directly into the learning process, resulting in a unified framework that integrates scientific machine learning (SciML) with data assimilation. The discrete-time CGKN is tested on canonical nonlinear PDE systems with intermittent and turbulent features, demonstrating forecasting performance on par with state-of-the-art SciML methods and delivering efficient, accurate DA with uncertainty quantification. Beyond DA, the CGKN framework exemplifies how SciML models can be designed to seamlessly interface with outer-loop applications such as design optimization, inverse problems, and

optimal control.

Zhongrui Wang
University of Wisconsin-Madison
Mathematics Department
wang3262@wisc.edu

Nan Chen, Nan Chen
University of Wisconsin - Madison
chennan@math.wisc.edu, chennan@math.wisc.edu

MS106

A Tractable Nonparametric Probabilistic Framework for Model-Form Uncertainty in Turbulent CFD

We present a computationally tractable approach for modeling and quantifying model-form uncertainty in viscous CFD. The method distinguishes between two sources: turbulence-modeling uncertainty and all other sources (e.g., boundary conditions). Accordingly, it combines two complementary strategies. The first targets turbulence modeling by decomposing the Reynolds stress tensor into deviatoric and spherical parts; constructing a hyperparameterized probabilistic model of the eigenvalues of the deviatoric component based on its spectral properties; and extending this to a complete probabilistic representation of the Reynolds stress. Each realization corresponds to an admissible turbulence model within a family. The second strategy adapts a recently developed nonparametric probabilistic method to this context. Grounded in stochastic, projection-based model order reduction (PMOR), it ensures computational tractability while capturing additional model-form and parametric uncertainties. Hyperparameters of both components are identified by minimizing a data-driven stochastic loss function, with the framework also accounting for the uncertainty introduced by PMOR itself to enable efficient Monte Carlo simulation. The approach is demonstrated on RANS-based aerodynamic analysis of the NASA Common Research Model in transonic flow, where wind tunnel data is available, highlighting its potential for large-scale CFD uncertainty quantification with turbulence modeling.

Charbel Farhat
Stanford University
cfarhat@stanford.edu

MS106

Embedded Model Error Using Orthogonal Gaussian Processes

In this work, we extend the embedded model error framework to incorporate flexible model error representations that distinguish uncertainy due to model error from that resulting from measurement noise. In an inverse modeling context, we calibrate a simplified model using data from a more detailed predictive model treated as the truth in the presence of measurement noise. Specifically, we embed a Gaussian process (GP) in the model-to-be-calibrated, enabling flexible model corrections. For tractability, we adopt the weight-space GP perspective, and we examine its equivalence to the function-space view in both interpolative and extrapolative scenarios through numerical experiments. To separate the predictions of the model from those of the embedded GP, we extend the orthogonal GP method, originally developed for external GPs, to the embedded setting. Specifically, we propose two strategies: (i) linearizing the embedded GP model, and (ii) introducing regularization in the GP prior to enforce orthogonality between the GP and the model gradients in the parameter space. Tests on synthetic linear and nonlinear examples demonstrate that the framework provides flexible corrections, with orthogonality constraints effectively decorrelating model parameters and GP weights, thereby enabling efficient posterior sampling of the model and its correction.

Mridula Kuppa
University of Illinois at Urbana-Champaign
vkuppa2@illinois.edu

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

Marco Panesi
University of California, Irvine
mpanesi@uci.edu

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS106

Analytic Solutions to Nonlinear Odes Via Spectral Power Series and Why Highly Reduced Models Can Work

A typical question in the field of MFE is: how can we improve a reduced set of ODEs or PDEs to better approximate a high-fidelity one? But we could also ask: how come these reduced models work as well as they do in the first place? In this talk, I'll present analytic solutions to nonlinear ODEs to explain why highly reduced models may work well, even for extrapolative scenarios. These solutions are found through a new method called spectral power series (SPS), which relies on eigenvalues of the linearized system around equilibrium. The exact coefficients of the series are found with linear algebra and some combinatorics. Finally, to ground SPS with existing dynamical systems theory, I'll show SPS as an explicit construction of the Koopman expansion and the Hartman-Grobman transformation.

Rebecca Morrison
CU Boulder
rebeccam@colorado.edu

MS106

Stochastic and Non-Local Closure Modeling for Nonlinear Dynamical Systems Via Latent Generative AI Methods

Closure models are widely used in simulating complex multiscale dynamical systems such as turbulence and Earths climate, for which direct numerical simulation that resolves all scales is often too expensive. For those systems without a clear scale separation, deterministic and local closure models often lack enough generalization capability, which limits their performance in many real-world applications. In this talk, we present a data-driven modeling framework for constructing stochastic and nonlocal closure models via latent conditional diffusion model and neural operator. We also compare the implicit regularization of the latent space offered by a joint training scheme against two explicit approaches that incorporate the latent space regularization, via metric-preserving (MP) or geometry-aware (GA) constraints. The results show that the proposed methodology

provides a systematic approach via latent generative AI methods to constructing efficient and generalizable data-driven stochastic closure models for multiscale dynamical systems with continuous spatiotemporal fields.

Xinghao Dong, Huchen Yang, Chuanqi Chen, Jinlong Wu
University of Wisconsin-Madison
xdong94@wisc.edu, huchen.yang@wisc.edu, cchen656@wisc.edu, jinlong.wu@wisc.edu

MS107

A General Framework for Convergence of Statistical Solutions and Applications

A trajectory statistical solution of a given evolution equation is a probability measure on the space of trajectories whose support is contained in the corresponding set of solutions, in addition to obeying a suitable regularity condition. I will present a general result yielding a set of verifiable assumptions that yield convergence of statistical solutions of a given parametrized family of evolution equations towards a statistical solution of a limiting system. As illustration, I will apply this result to the inviscid limit of the Navier-Stokes equations towards the Euler equations, and to Galerkin approximations of the Navier-Stokes equations.

Cecilia Mondaini
Drexel University
cf823@drexel.edu

MS107

Structure-Aware Warm-Starts for Uncertainty Propagation in Transport-Dominated PDEs

Transport-dominated phenomena pose difficulties for physics-informed machine learning (PIML), because training of PIML models can become unstable or inaccurate due to stiffness in the optimization landscape, spectral bias toward smooth solutions, and strong sensitivity to initialization and loss balancing. These challenges impact uncertainty quantification (UQ) as it demands high-fidelity forward solves across a high-dimensional uncertainty space, while surrogate models can be brittle; especially as advection speed increases and the problem becomes more ill-conditioned for gradient-based training. We present a structure-aware warm-start framework for efficient UQ in transport-dominated PDEs based on numerical continuation and transfer learning to stabilize and accelerate training of surrogate models. The core idea is to construct a sequence of regularized problems (e.g., reduced advection strength, increased regularization, coarser discretization) and to warm start each stage using the solution from the previous stage. This continuation path is combined with transfer across related geometries or dimensions (e.g., from 1D to 2D advection) to reduce the marginal cost of additional samples and enable rapid exploration of parameter regimes. We demonstrate that the proposed warm-start strategy improves robustness and reduces training time for uncertainty propagation compared with naive random-start baselines, while maintaining accuracy in solution profiles. The framework is compatible with Monte Carlo and collocation-style sampling, and it provides a practical route to scalable UQ when high-fidelity solves of transport-dominated PDEs are the computational bottleneck.

Daniel M. Tartakovsky
Stanford University
tartakovsky@stanford.edu

Hongli Zhao
University of Chicago
Committee on Computational and Applied Mathematics
honglizhaobob@uchicago.edu

MS108

Quantifying Uncertainty in Grid Topology Mappings: A Distribution Free Risk Control Approach

Accurate recovery of grid topology from overhead geospatial imagery underpins simulation, reliability, and resilience analysis of the electricity grid infrastructure. Errors in computer-vision outputs – missed towers, spurious detections – propagate to downstream tasks and can misinform planning decisions. We develop a distribution-free risk-control framework for uncertainty quantification of grid-topology extraction that allows us to separate existence from geometric quality of model predictions. Specifically we define existence-centric risks (recall shortfall and false-discovery proportion) for detection that go beyond IoU based metrics, and use Conformal Risk Control (CRC) and Learn-then-Test (LTT) frameworks to provide explicit, finite-sample statistical guarantees. Our procedure is detector-agnostic, yields interpretable per-detection probabilities for downstream tasks, and can be merged into existing pipelines. It directly supports risk-aware decisions in resilience, reliability, and investment planning for power-grid infrastructure. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Document release number: LLNL-ABS-2011544.

Jayanth Jagalur
Lawrence Livermore National Laboratory
jagalur1@llnl.gov

Jaffe Luke, Juliette Franzman
LLNL
jaffe5@llnl.gov, franzman1@llnl.gov

Christina Mao, Mason Blake, Jean-Paul Watson
Lawrence Livermore National Laboratory
mao6@llnl.gov, mason35@llnl.gov, watson61@llnl.gov

MS108

Scalable Transport Maps to Enable Large-Scale Uncertainty Quantification for NASA Remote Sensing Missions

Earth-observing satellite missions are capturing information at increasingly fine spatial and temporal scales, providing great potential for science, applications, and informed decision making. However, the corresponding massive data volumes place increasing demands on computational resources and limit feasible implementation of standard uncertainty quantification (UQ) practices, e.g. Monte Carlo. Additionally, the retrieval of high priority geophysical variables, like evapotranspiration (ET), relies on algorithms that ingest variables from multiple sources with varying fidelity of UQ. These inputs are typically spatial fields with uncertainties that exhibit complex spatial structure that must be propagated to ensure reliable decision-making from field-level to global scales. In this talk, we present a novel implementation of transport maps to enable computationally feasible, forward UQ for ET products from NASAs ECOSTRESS mission. We propose a hierarchical transport map that is flexible to capture arbitrary input uncertainty distributions varying across space and

time. The resulting map amortizes the cost of performing UQ across a range of geophysical conditions (e.g. multiple ecosystems), thereby enabling scalable UQ for operational satellite data pipelines. We demonstrate the methodology over agricultural regions with complex space-time dynamics, and we discuss ongoing work to extend the methods to capture spatially correlated uncertainty information.

Maggie Johnson, Jonathan Hobbs
Jet Propulsion Laboratory
maggie.johnson@jpl.nasa.gov,
jonathan.m.hobbs@jpl.nasa.gov

Ricardo Baptista
University of Toronto
r.baptista@utoronto.ca

Gregory Halverson
Jet Propulsion Laboratory
gregory.h.halverson@jpl.nasa.gov

Kerry Cawse-Nicholson
Jet Propulsion Laboratory
California Institute of Technology
kerry-anne.cawse-nicholson@jpl.nasa.gov

Otto Lamminpää
Jet Propulsion Laboratory
otto.m.lamminpaa@jpl.nasa.gov

Meredith Franklin
University of Toronto
meredith.franklin@utoronto.ca

MS108

Validation of UQ Methods for AI Surrogate Models

While it is essential to characterize the trustability of AI models, uncertainty quantification remains an open question for this field. Even for the narrow scope of AI surrogates with numeric inputs and outputs, there are many different approaches to quantifying uncertainty, but these methods do not agree. In this talk, we discuss our efforts to evaluate the performance of these differing UQ methods. LLNL-ABS-2011553. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Kathleen Schmidt, Vivek Narayanaswamy, Kowshik Thopalli, Boya Zhang
Lawrence Livermore National Laboratory
schmidt41@llnl.gov, narayanaswam1@llnl.gov,
thopalli1@llnl.gov, zhang67@llnl.gov

Amanda Muyskens
Real Good AI
almuyskens@realgoodai.org

MS108

Deep Gaussian Process Regression for Shock Modeling

Accurate modeling of shock phenomena is critical in the field of materials science, where predicting responses under extreme conditions is essential. Statistical emulators of high fidelity simulations make it possible to perform the large number of evaluations necessary for uncertainty quantification. However, traditional Gaussian process (GP) approaches often struggle with the inherent nonstationarity of shock data and other ML methods lack probabilistic interpretability. Instead, Deep Gaussian Process Regression (DGPR) is a powerful, flexible framework for emulation that extends standard Gaussian processes to accommodate complex response behavior by stacking multiple layers. We use a DGP to emulate a shock response and calibrate to real-world shock data, comparing performance to a standard GP.

Laura Wendelberger
Lawrence Livermore National Laboratory
wendelberger1@llnl.gov

MS109

Diffusion Annealed Langevin Monte Carlo for Generative Modelling and Sampling

In this talk, we study the theoretical properties of general diffusion paths and their Langevin Monte Carlo implementation, termed diffusion annealed Langevin Monte Carlo (DALMC), under weak assumptions on the data distribution. We derive non-asymptotic error bounds for annealed Langevin dynamics where the path of distributions is given by Gaussian convolutions of the data, as in diffusion models. We then extend the analysis to heavy-tailed (Students t) diffusion paths, establishingfor the first timetheir theoretical validity for heavy-tailed data. These results provide rigorous guarantees for a class of score-based generative models that interpolate between a simple distribution (Gaussian or Students t) and the data in finite time, offering a broader perspective than standard diffusion approaches based on a forward OrnsteinUhlenbeck process. Building on this framework, we introduce new methods for efficient sampling from complex, unnormalised targets via multiscale dynamics. By exploiting stochastic averaging in a slowfast system of stochastic differential equations, we estimate intermediate scores along the diffusion path without training or inner-loop MCMC. We present two algorithms: MultALMC, based on multiscale annealed Langevin dynamics, and MultCDiff, which uses multiscale controlled diffusions for the reverse-time OU process. Both overdamped and underdamped variants are developed, with theoretical convergence guarantees to the target diffusion path.

Paula Cordero-Encinar, Deniz Akyildiz
Imperial College London
paula.cordero-encinar22@imperial.ac.uk,
deniz.akyildiz@imperial.ac.uk

Sebastian Reich
Universität Potsdam, Germany
sereich@uni-potsdam.de

Andrew Duncan
Imperial College London
a.duncan@imperial.ac.uk

MS109

Preconditioning for Mcmc: Concepts, Practices, and Challenges

Motivated by the improvement of sampling algorithms, the constant creation of preconditioners leaves a slew of new techniques in the wake of each passing year [Titsias 2023, Cui et al. 2024, Cai et al. 2024]. Despite this, the notion of preconditioning in the minds of practitioners and theo-

reticians is vague and incomplete. Each new preconditioner comes with its own justification, particular to each individual problem setting. In the talk I will formalise the notion of preconditioning, unifying existing and future preconditioners under a single conceptual framework. I will then lead the audience through a catalogue of potential hurdles to the good operation of a sampler. In each case I will explain why the hurdle obstructs the sampling operation and I will introduce and test the preconditioners that have been posed to surmount them. Ill provide either theory or empirical examples to verify claims as to the malignance of the obstacle and the effectiveness and behaviour of the preconditioner. Then I will introduce a taxonomy for preconditioners: those that modify the target distribution, and those that modify the algorithm. I will describe the natural capabilities of the preconditioners that fall into these classes. Lastly I will introduce an important concept informing the construction of preconditioners, which we call 'Gaussian Equivalence. Joint work with Sam Power

Max Hird
University of Waterloo
max@hirds.co.uk

MS109

Optimizing Irreversible Perturbations of the Unadjusted Langevin Algorithm

In this talk, we investigate irreversible perturbations of overdamped Langevin dynamics as a strategy to accelerate sampling. These perturbations add skew-symmetric drift terms that preserve the target distribution while inducing rotational flows, thereby improving exploration. While their advantages are well understood in continuous time, the role of discretization and the design of optimal perturbations for discrete-time algorithms remain largely unexplored. Focusing on the Unadjusted Langevin Algorithm (ULA), we develop a framework for optimizing constant irreversible perturbations that balances two competing objectives: maximizing the spectral gap of the continuous dynamics to speed convergence, and minimizing discretization error that induces bias. We introduce new criteria that jointly quantify bias and efficiency, and show how they identify perturbations that improve sampling performance beyond existing constructions.

Julie Zhu
MIT
qianyu_z@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

Konstantinos Spiliopoulos
Boston University
Department of Mathematics and Statistics
kspiliop@bu.edu

Benjamin Zhang
University of North Carolina at Chapel Hill
bjz@unc.edu

MS110

Extreme Scale Bayesian Inference and Optimal Experimental Design for Wave Propagation-based Source Inversion

We consider Bayesian source inversion problems and associated optimal experimental design governed by wave propagation in very high dimensional parameter space. This class of problems is fundamentally challenging since it does not lend itself to low dimensional approximation. We achieve a real-time Bayesian inversion capability by exploiting the time-shift invariance of the wave propagation problem and resulting parameter-to-observable map, an offline–online decomposition with a modest number of offline (adjoint) wave propagation solutions along with a wave propagation-free online component, and novel parallel algorithms for the online component that map well onto GPU clusters. Capitalizing on this capability, we design novel greedy algorithms for solving sensor selection OED problems with the goal of maximizing expected information gain for both the parameter field as well as posterior predictive quantities of interest. The real-time Bayesian inversion and OED framework is demonstrated for acoustic–gravity wave propagation-based tsunami early warning in several high-risk subduction zones settings. Bayesian inverse problems with up to 1 billion inversion parameters representing spatiotemporal seafloor motion are solved in less than a second on 512 GPUs, and OED sensor selection problems for tsunami early warning with hundreds of acoustic pressure sensors are solved to near optimality in a few hours.

Sreeram R. Venkat
University of Texas Austin
201 E 24th St, Austin, TX 78712
srvenkat@utexas.edu

Stefan Henneking
University of Texas at Austin
stefan@oden.utexas.edu

Alice Gabriel
University of California, San Diego
algabriel@ucsd.edu

Omar Ghattas
University of Texas at Austin
omar@oden.utexas.edu

MS110

Shock-Capturing Data Assimilation Via Hierarchical Bayes, EnKF, and Discontinuous Galerkin

Data assimilation methods are currently used to tackle many practical, applied problems throughout computational science. At their heart, they gauge the state of a targeted dynamical system, and many workhorse methods (EnKF, ETKF, etc.) use ensembles to represent the uncertainty over this target state. We approach these methods with a particular focus toward hyperbolic conservation laws, giving piecewise-smooth states with possibly large discontinuities. While many works focus on representing the uncertainty over the dynamical system state more rigorously by using nonlinear transformations, they are often agnostic of the knowledge over the dynamical system they quantify uncertainty for. Contrasting with the (localized) ensemble Kalman filter, we demonstrate the ability to ensure that each *sample* drawn from our data assimilation method is a realistic realization of the underlying partial differential equations (PDEs), and further extend our method to take advantage of underlying structure in discontinuous Galerkin formulations of the PDEs.

Daniel Sharp, Jan Glaubitz, Mathieu Le Provost, Youssef M. Marzouk

Massachusetts Institute of Technology
dannys4@mit.edu, Glaubitz@mit.edu, n/a, ymarz@mit.edu

MS110

Hierarchical Prior-Informed Dictionary Learning Methods

Dictionary matching methods seek to match a given data vector to measured or synthetically generated dictionary templates that can be interpreted to constitute software-informed priors. The problem can be recast as a sparse coding problem. To facilitate the process, it is often useful to parcel the dictionary into subdictionaries using some coherence criterion, and then identify the relevant subdictionaries leading to a deflated coding problem. In this talk, we review Bayesian methods that improve the identifiability of the relevant subdictionaries and lead to improved performance and interpretability, and lower computational cost than solving the original matching problem. The methodology is demonstrated with applications from different areas, e.g., brain imaging.

Erkki Somersalo
Case Western Reserve University
ejs49@case.edu

MS110

Using Sciml to Learn a Prior from Population Data

Stochastic inverse problems have been attracting increasing attention in recent years due to advances in data acquisition techniques which enable utilizing tremendous amounts of data to construct data-informed physics-based computational models. In particular, a recently developed approach, called data-consistent inversion (DCI), solves a specific class of stochastic inverse problems where one seeks an aleatoric characterization of model inputs over a population of assets/individuals. Moreover, this approach has been recently utilized to construct population-informed priors to enhance the information gained in performing asset-specific Bayesian inference for digital twins. At the same time, scientific machine learning methods have also become much more prevalent in computational science, largely due to their ability to learn and exploit low-dimensional structure in high-dimensional data. The subject of utilizing approximate models/maps in the context of DCI has been studied in extensive detail. In this presentation, we will describe some of our recent work on using data-driven (SciML) surrogate models to expedite the solution to both Bayesian and data-consistent inverse problems.

Tim Wildey
Optimization and Uncertainty Quantification Dept.
Sandia National Laboratories
tmwilde@sandia.gov

MS111

Marginalization-Based Regularization in Separable Nonlinear Bayesian Inverse Problems

Separable nonlinear inverse problems arise in many applications where a forward model depends linearly on some unknowns and nonlinearly on others, with semi-blind deconvolution as a motivating example. In this talk, we adopt a Bayesian perspective with Gaussian observation noise and Gaussian priors on the linear variables, and show that common choices of priors for the nonlinear parameters lead to maximum a posteriori estimates equivalent to regularized separable nonlinear least-squares problems, which can be solved efficiently via variable projection. We then show that analytically marginalizing the linear variables yields a reduced posterior for the nonlinear parameters whose objective augments the usual variable-projection residual with a principled log-determinant term: this term encodes uncertainty in the eliminated variables, discourages degenerate no-blur solutions, and admits scalable evaluation for structured forward models. Numerical results in semi-blind deconvolution illustrate improved stability and reconstruction quality of the proposed marginalization-based regularization.

Malena I. Espanol, Jordan T. Dworaczyk
Arizona State University
malena.espanol@asu.edu, jtdworac@asu.edu

MS111

Priorconditioned Sparsity-Promoting Projection Methods for Deterministic and Bayesian Linear Inverse Problems

Rapidly-growing fields like data science and machine learning require fast, accurate methods for inverse problems. Key challenges include ill-posedness, high dimensionality, and complex model constraints. Sparsity-promoting techniques for deterministic and Bayesian regularization address these by solving sequences of high-dimensional iteratively reweighted least squares (IRLS) problems on lower-dimensional subspaces. Generalized Krylov subspace (GKS) methods efficiently solve these IRLS sequences by projecting large-scale problems into a small, successively enlarged subspace. However, these sparsity-promoting GKS (S-GKS) methods converge slowly. We address this by combining S-GKS with priorconditioning (PS-GKS). Integrated into the iterative alternating sequential (IAS) algorithm, PS-GKS enables automatic selection of the generalized gamma hyper-prior's shape/rate parameter. We also explore restarting and recycling variants that handle memory constraints. Numerical experiments confirm that PS-GKS and its variants are competitive with or outperform existing hybrid Krylov methods.

Mirjeta Pasha
Virginia Tech
mpasha@vt.edu

MS111

Probabilistic Imaging Using Diffusion Image Priors

Probabilistic imaging frames reconstruction as posterior inference, leveraging a physical forward model (likelihood) and an expressive image prior for both high-quality reconstructions and uncertainty quantification. In this talk, we extend diffusion image priors beyond the standard known-operator setting to blind inverse problems, where key parts of the measurement model are unknown and must be inferred jointly with the image. We present PRISM, a diffusion-based Bayesian framework that places diffusion priors on the image and on the unknown operator parameters, and performs posterior sampling via a principled alternating (Gibbs-style) scheme. A central ingredient is a measurement-conditioned diffusion model for the operator, which injects observation information to improve robustness under poor initialization and model mismatch. We demonstrate PRISM on blind motion deblurring, showing improved recovery and credible uncertainty estimates over

recent diffusion baselines.

Yu Sun
Johns Hopkins University
ysun214@jh.edu

MS111

Efficient Simulation of High-Resolution Turbulent Systems Using Diffusion Models

Simulating large and high-resolution turbulent systems can be computationally prohibitive. While modern machine learning models such as diffusion models offer promising generative capabilities, generating high-resolution turbulent fields remains challenging due to GPU memory limitations arising from the large input size. In this work, we address this challenge by proposing a hierarchical strategy based on diffusion models. We first train a low-resolution model to generate the global dynamics of turbulence. In parallel, we train patch-wise diffusion models to predict high-resolution local patches, conditioned on previous high-resolution patches and the current low-resolution global field. These patches are then combined using a partition of unity to produce a seamless high-resolution solution. In our framework, each model only processes a small portion of the domaineither the low-resolution global field or localized high-resolution patchessignificantly reducing memory requirements compared to end-to-end training at full resolution. Moreover, the method is naturally parallelizable by design. Numerical experiments demonstrate the performance and efficiency of the proposed approach.

Zezhong Zhang
Auburn University
zez0002@auburn.edu

MS112

Sparse Generative Ml Driven Autonomous Solps Grid Generation.

Automatic mesh generation for boundary modeling codes that adjusts to scans in geometry, physical constraints is an essential part of computational aided design. We have developed a capability to optimize SOPs grids given physically relevant local or global loss functions. This talk will describe an AI/ML framework were grids produced by experts or optimization can be used to train generative models. Framework is adaptable, capable of learning new reactor designs and responding to community input.

Rick Archibald
Data Analytics and Machine Learning Group
Oak Ridge National Labortory
archibaldrk@ornl.gov

MS112

MaLT: Machine-Learning-Guided Test Case Design and Fault Localization of Complex Software Systems

Software testing is essential for the reliable and robust development of complex software systems. However, testing and fault localization of such systems can be very costly. To mitigate this cost, we outline in this work a holistic machine-learning-guided test case design and fault localization (MaLT) framework, which leverages recent probabilistic machine learning methods to accelerate the testing of complex systems. MaLT consists of three steps: (i) the construction of a suite of test cases using a covering array for initial testing, (ii) the investigation of posterior root cause probabilities via a Bayesian fault localization procedure, then (iii) the use of such Bayesian analysis to guide selection of subsequent test cases via active learning. The proposed MaLT framework can thus facilitate efficient identification and subsequent diagnosis of software faults with limited test runs.

Irene Ji
JMP Statistical Discovery LLC
irene.ji@jmp.com

Simon Mak
Duke University
sm769@duke.edu

Ryan Lekivetz, Joseph Morgan
JMP Statistical Discovery LLC
ryan.lekivetz@jmp.com, joseph.morgan@jmp.com

MS112

VVUQ for Reduced Models of Edge Transport Barriers in Tokamak Fusion Devices

The performance of tokamak fusion devices in high confinement mode depends on an edge transport barrier which increases the achievable plasma pressure. However, predictions of plasma temperature and density in the transport barrier remain limited by gaps in our understanding of turbulent transport. While high-fidelity nonlinear simulations are computationally expensive and impractical for routine use, existing reduced models rely on transport constraints that limit their applicability in relevant plasma scenarios. This work addresses these limitations by developing fast, validated quasilinear transport models for the pedestal. The models are based on a mixing-length approach, utilizing linear gyrokinetic simulations performed with GENE. A Bayesian uncertainty quantification framework is employed to calibrate and validate the models against experimental pedestal profiles from the DIII-D tokamak, incorporating experimental uncertainties via Gaussian process regression. The forward propagation of uncertainties is implemented using the integrated modeling framework ASTRA. Initial validation results are presented for two transport models: one based on electron temperature gradient (ETG) modes [Hatch et al., Nucl. Fusion, 2024] and another targeting electromagnetic modes.

Leonhard Leppin
University of Austin
leppin@utexas.edu

Cole Stephens
University of Texas at Austin
cole.stephens@austin.utexas.edu

Pieterjan Robbe
Sandia National Laboratories
pmrobbe@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Ping-Yu Li, Joseph Schmidt, Saeid Houshmandyar
University of Texas at Austin
pingyu.li@austin.utexas.edu, joseph.schmidt@utexas.edu,
houshmandyar@austin.utexas.edu

Norman M. Cao
Institute for Fusion Studies
The University of Texas at Austin
norman.cao@austin.utexas.edu

Caitlin Curry
Sandia National Labs
cjcurry@sandia.gov

David Hatch
University of Texas at Austin
drhatch@austin.utexas.edu

Todd Oliver
The University Of Texas-Austin
oliver@oden.utexas.edu

MS112

Evidence Aggregation from VVUQ Activities Towards Quantitative Model Credibility Assessment

A computational approach is developed to aggregate information from a variety of model assessment activities and provide the decision-maker with quantitative information relating to model credibility. Model credibility is defined here as the degree to which the decision-maker believes that using the model prediction is acceptable for its intended application, based on the collected evidence. The proposed credibility measure makes use of six assessment activities considered in the Predictive Capability Maturity Model (PCMM) developed at Sandia National Laboratories. Based on the maturity, strength and relevance of evidence provided by the six PCMM elements, Dempster-Shafer theory (DST) is employed to combine the assessment scores and provide a three-dimensional metric of model credibility: a [0,1] measure of the degree of evidentiary support for the model's acceptability for its intended use, a [0,1] measure of the degree of evidentiary support against the model's acceptability for its intended use, and a [0,1] measure of the uncertainty about a model's acceptability due to lack of mature and/or relevant evidence. Extension of this credibility metric to assess model risk, i.e., risk of making wrong decisions, is discussed. This work was funded by Sandia National Laboratories.

Sankaran Mahadevan
Vanderbilt University
sankaran.mahadevan@vanderbilt.edu

Oliver T. Stover
Vanderbilt University
Civil Engineering
oliver.stover@vanderbilt.edu

Pranav Karve
vanderbilt
pranav.m.karve@vanderbilt.edu

Kyle Neal, Joshua Mullins
Sandia National Laboratories
kneal@sandia.gov, jmullin@sandia.gov

MS113

Taylor Series Methods for Aerocapture Uncertainty Quantification

This paper presents an analytical approach for uncertainty quantification in aerocapture scenarios using higher-order augmented directional state transition tensors with Gaussian mixture models to efficiently propagate uncertainties and represent nonlinear state distributions. State transition tensors (STTs), higher-order Taylor series expansions of the solution flow, have been widely used as a computationally efficient analytical uncertainty quantification method for orbital scenarios, but have not previously been applied to aerocapture. However, obtaining higher-order STTs requires integrating exponentially more equations. Directional state transition tensors (DSTTs) mitigate this cost by projecting the state into a reduced-dimension basis. This work develops new dynamics analysis techniques to identify effective bases for this reduction, including decomposed and higher-order Cauchy Green tensors and augmented formulations tailored to quantities of interest such as apoapsis radius. These dynamics-informed directions are used for state transition tensor basis reduction and Gaussian mixture model mixand splitting. The proposed method is demonstrated in a Uranus aerocapture scenario, with results compared against Monte Carlo simulations and the Unscented Transform.

Grace Calkins
University of Colorado Boulder
grace.calkins@colorado.edu

David Woffinden
NASA Johnson Space Center
david.c.woffinden@nasa.gov

Jay McMahon
University of Colorado Boulder
jay.mcmahon@colorado.edu

MS113

A Parallel Grid-based Algorithm for Orbital Uncertainty Propagation

Eulerian nonlinear uncertainty propagation methods often suffer from finite domain limitations and computational inefficiencies. A recent approach to this class of algorithm, Grid-based Bayesian Estimation Exploiting Sparsity, addresses the first challenge by dynamically allocating a discretized grid in regions of phase space where probability is non-negligible. However, the design of the original algorithm causes the second challenge to persist in high-dimensional systems. This paper presents an architectural optimization of the algorithm for CPU implementation, followed by its adaptation to the CUDA framework for single GPU execution. The algorithm is validated for accuracy and convergence, with performance evaluated across distinct GPUs. Tests include propagating a three-dimensional probability distribution subject to the Lorenz '63 model and a six-dimensional probability distribution subject to the Lorenz '96 model. The results imply that the improvements made result in a speedup of over 1000 times compared to the original implementation. We leverage this speedup to feasibly propagate the six-dimensional state uncertainty of a highly inclined Saturn-Enceladus halo orbit in the Circular Restricted Three-Body Problem.

Benjamin L. Hanson
University of California San Diego
blhanson@ucsd.edu

Carlos Rubio, Adrián García-Gutiérrez
University of León
crubs@unileon.es, a.gutierrez@unileon.es

Thomas Bewley, Aaron Rosengren
University of California San Diego
tbewley@ucsd.edu, arosengren@ucsd.edu

MS113

Conjunction Analysis in Multi-Body Mechanical Systems

The problem of quantifying the probability of collision between two rigid bodies that exhibit general relative motion is considered. Traditional approaches to evaluate the collision probabilities between spheres in three degrees of freedom are discussed with special consideration on long-term encounters between the particles of interest. Engineering approximations are presented that enable the evaluation of the collision probabilities between rigid bodies as a sum of the particle collision probabilities. Torque free motion of a rigid body is used to enable the semi-analytical evaluation of the relative motion integrals involved in the conjunction assessment. Applications of the technique to evaluate the probability of collision of an articulating space object performing servicing maneuvers with respect to a client vehicle are used to demonstrate the utility of the methodologies developed in this paper. Numerical methods for efficient evaluation of the collision probabilities are discussed. It is shown that the technical approaches to evaluate the collision probability in multibody mechanical systems are also useful in vehicle jettison analysis in astronautics.

Manoranjan Majji, Luke Murphy, Thomas Ahrens
Texas A&M University
mmajji@tamu.edu, lukeryan5000@tamu.edu, tahrens@tamu.edu

MS114

From Single to Multiple Kernels: Learning Interacting Particle Systems on Networks

We present a framework for jointly learning the network structure, interaction kernels, and type assignments in heterogeneous interacting particle systems (IPS) on networks. The agents follow first-order stochastic dynamics, with interactions governed by unknown kernel functions depending on discrete type pairs. Since species information is not observed a priori, the learning problem is inherently a nonconvex mixed-integer optimization. To address this challenge, we decouple the task into a low-rank matrix sensing stage and a clustering stage, followed by decomposition and post-processing to refine estimates. The matrix sensing step leverages shared kernel structure to recover a low-rank embedding of network and kernel parameters, while the clustering step infers discrete types via angular separation of normalized embeddings. We establish theoretical guarantees under restricted isometry properties (RIP), and validate the method on synthetic examples, including heterogeneous predatorprey systems. Our approach generalizes to a broad class of multitype IPS models and provides a scalable, robust solution for learning structured particle interactions from trajectory data.

Quanjun Lang
Duke University
quanjun.lang@duke.edu

MS115

Data-Driven MoriZwanzig with Variable Projection Operators

The MoriZwanzig (MZ) formalism decomposes an observable into resolved and unresolved components through a projection operator, and the accuracy of the resulting reduced-order model is critically dependent on this choice. Traditionally, the projection operator is fixed and applied recursively to derive the memory terms, which capture the intrinsic non-Markovian character of the MZ formalism. In earlier work, we demonstrated that the projection operator can be defined implicitly as a learning problem, specified by a model class and an empirical risk function to be minimized. However, fixing the same projection operator imposes a constraint: both the model class and cost function must remain identical across the Markov term and all memory operators. In this work, we relax this restriction by allowing variable projection operators within the MZ framework. This generalization enables a data-driven approach in which different components of the MZ expansion can be modeled flexibly with mixed models. We present theoretical insights into this extended formulation and provide numerical results illustrating that the variable-projection MZ framework yields more accurate and adaptable reduced-order models compared to the conventional fixed-operator approach.

Yen Ting Lin
Los Alamos National Laboratory
yentingl@lanl.gov

MS115

Differentiable Dynamic Mode Decomposition

We propose Differentiable Dynamic Mode Decomposition (Diff-DMD), a framework that enables gradient-based design and analysis of nonlinear dynamical systems directly through functions of DMD eigenvalues and modes. By embedding an adjoint-based sensitivity formulation into the standard DMD pipeline, Diff-DMD computes exact derivatives of modal quantities with respect to system parameters and data at the cost of a single forward factorization plus one adjoint solveindependent of the number of parameters. We derive sensitivities for both independently distributed and sequential snapshot constructions and further obtain data sensitivities to support robust design and adversarial data analysis. The approach yields orders-of-magnitude speedups over finite-difference approximations while maintaining accuracy. We demonstrate the method on (i) stabilizing a linear system by shifting the real parts of its eigenvalues, (ii) suppressing chaos in the Lorenz system by steering parameters until all continuous-time DMD eigenvalues have negative real parts, and (iii) increasing the oscillation frequency of a seven-dimensional yeast glycolysis model under stability and slow-decay regularization. Across these cases, Diff-DMD produces stable, accurate gradients that agree with finite differences and enable efficient optimization in high-dimensional settings, offering a scalable surrogate for design, control, and system identification based on Koopman/DMD modal structure.

Shaowu Pan
RPI
pans2@rpi.edu

MS115

Structure-preserving Lift and Learn: Scientific Machine Learning for Nonlinear Conservative Partial

Differential Equations

This talk presents structure-preserving Lift & Learn, a scientific machine learning method that employs lifting variable transformations to learn structure-preserving reduced-order models for nonlinear partial differential equations (PDEs) with conservation laws. We propose a hybrid learning approach based on a recently developed energy-quadratization strategy that uses knowledge of the nonlinearity at the PDE level to derive an equivalent quadratic lifted system with quadratic system energy. The lifted dynamics obtained via energy quadratization are linear in the old variables, making model learning very effective in the lifted setting. Based on the lifted quadratic PDE model form, the proposed method derives quadratic reduced terms analytically and then uses those derived terms to formulate a constrained optimization problem to learn the remaining linear reduced operators in a structure-preserving way. We demonstrate the generalizability of quadratic models learned via the proposed structure-preserving Lift & Learn method through three numerical examples: the one-dimensional wave equation with exponential nonlinearity, the two-dimensional sine-Gordon equation, and the two-dimensional Klein-Gordon-Zakharov equations. The numerical results show that the proposed learning approach is competitive with the state-of-the-art structure-preserving data-driven model reduction method in terms of both accuracy and computational efficiency.

Harsh Sharma
University of WisconsinMadison
hsharma33@wisc.edu

MS116

Information Theory Based Active Learning For High Dimensional Regression and Operator Learning

Deep learning has found extensive success for generative and classification tasks. This success is partially due to the large amounts of data available for training large neural network models. In scientific computing, however, training sets can be scarce and new data expensive to collected, often requiring long numerical simulations or even physical experiments that cannot be done in-silico. Efficient data collection has the potential to help increase accuracy and model robustness in this setting. Active learning takes into consideration model (or epistemic) uncertainty in order to sample data points that have the largest potential of decreasing model uncertainty. In this project, we present a framework for active learning in regression and operator learning, practically extending the popular acquisition function BALD to high-dimensional regression. Our methodology is defined using epistemic neural networks, which provide a general framework that applies to Bayesian Neural Networks, Deep Ensembles, MC Dropout and more. By running numerical experiments, we find a clear advantage to using our active learning methods over variance sampling.

Leonardo Ferreira Guilhoto, Akshat Kaushal, Paris Perdikaris
University of Pennsylvania
guilhoto@sas.upenn.edu, akaush@seas.upenn.edu, pgp@seas.upenn.edu

MS116

Active Operator Learning over Shapes and Other Function Spaces

Active learning has proven successful for robust training of surrogates in settings where data generation can be controlled. Such methods typically rely on uncertainty-aware surrogates to drive efficient dataset generation, and as a result are most established for learning functions that map between vector spaces. In recent years, there has been increasing interest in non-vector structured data, like graph- and field-structured data, as they are common in the scientific domain (e.g. spatio-temporal processes). To that end, we explore ways to quantify the uncertainty of such operator approximations, particularly for operators that map shapes to quantities of interest. Using functional or graph shape representations, we develop such operators for closed surfaces with UQ. Secondly, we explore ways to generate a pool of valid candidates (e.g. shapes) for active learning, which is not as trivial as generating candidates in a vector space. We show how these two ingredients can be coupled to yield an effective active learning recipe for operators.

Michalis Michaelides
PhysicsX
michalis.michaelides@physicsx.ai

MS116

Active Learning for Neural PDE Solvers

Efficiently solving partial differential equations (PDEs) is critical across science and engineering. Neural PDE solvers have emerged as a promising alternative to established numerical solvers, offering fast and differentiable surrogate models. However, a main obstacle for neural PDE solvers is the high initial cost of obtaining the training dataset. In this talk, we will discuss active learning (AL) as a method to reduce data generation costs. By iteratively selecting only the most informative inputs to the simulator, AL can reach the same level of accuracy with fewer training samples. Yet, PDEs also pose a significant challenge for deep AL, as they involve an autoregressive, high-dimensional regression problem that differs fundamentally from the standard applications, such as image classification. To evaluate common deep AL strategies in the context of neural PDE solvers, we recently introduced the AL4PDE benchmark, which we describe in the second part of the talk. AL4PDE is a modular framework designed to accelerate the development of new AL strategies for neural PDE solvers. It contains ready-to-use neural PDE solvers, AL strategies, and PDE tasks. Based on AL4PDE, we present a benchmark study showing that deep AL methods can already reduce the average error compared to a randomly sampled dataset. In particular, AL achieves lower worst-case errors, making the final surrogate model more robust. Finally, we discuss recent advances in AL for PDEs.

Daniel Musekamp, Marimuthu Kalimuthu
University of Stuttgart
daniel.musekamp@ki.uni-stuttgart.de,
marimuthu.kalimuthu@ki.uni-stuttgart.de

David Holzmüller
INRIA
david.holzmuller@inria.fr

Makoto Takamoto
NEC
makoto.takamoto@neclab.eu

Mathias Niepert
University of Stuttgart

mathias.niepert@ki.uni-stuttgart.de

MS117

Continuum Transformers Perform In-Context Learning by Operator Gradient Descent

Transformers robustly exhibit the ability to perform in-context learning, whereby their predictive accuracy on a task can increase not by parameter updates but merely with the placement of training samples in their context windows. Recent works have shown that transformers achieve this by implementing gradient descent in their forward passes. Such results, however, are restricted to standard transformer architectures, which handle finite-dimensional inputs. In the space of PDE surrogate modeling, a generalization of transformers to handle infinite-dimensional function inputs, known as "continuum transformers,' has been proposed and similarly observed to exhibit in-context learning. Despite impressive empirical performance, such in-context learning has yet to be theoretically characterized. We herein demonstrate that continuum transformers perform in-context operator learning by performing gradient descent in an operator RKHS. We demonstrate this using novel proof strategies that leverage a generalized Representer Theorem for Hilbert spaces and gradient flows over the space of functionals on a Hilbert space. We further show the operator learned in context is the Bayes Optimal Predictor in the infinite depth limit of the transformer. We then provide empirical validations of this result and demonstrate that the parameters under which such gradient descent is performed are recovered through pre-training.

Abhiti Mishra
University of Michigan
abhiti@umich.edu

MS118

Towards a Multimodal Foundation Model for Heterogeneous Catalysis

The discovery of new catalysts is a major challenge, because it is limited by iterative, trial-and-error experimentation and computationally intensive simulations. Machine learning has shown promise to accelerate catalysis, but existing models are often restricted to a single data modality, unable to bridge the gap between vast textual knowledge in scientific literature and the explicit 3D atomic geometries of materials. To address this challenge, we introduce Alchemic, a natively multimodal foundation model that unifies natural language and 3D atomic structures. Our approach utilizes a Transformer backbone to learn the joint distribution over both data types. We achieve this through training a continuous tokenizer to embed structures in latent space; we then apply a flow matching objective to generate latent structures conditioned on the text input. Training with multimodality remains a significant challenge in AI-driven materials discovery and AI for science in general.

Daniel Israel
University of California Los Angeles
disrael@g.ucla.edu

MS118

Constrained Flow Matching for Scientific Applications

Generative models that satisfy hard constraints are essential in real-world applications, where physical laws or task requirements must be strictly enforced. Some approaches exist, but they struggle to guarantee feasibility. While repeated projection enforces feasibility, it distorts the learned distribution, leading to a mismatch with the data manifold. To mitigate this issue, existing methods introduce complex multi-stage procedures that apply projection to clean samples during sampling. However, these strategies typically increase complexity and suffer from error accumulation across steps. We seek to address these challenges by proposing a novel training-free method, Chance-constrained Flow Matching (CCFM). CCFM integrates stochastic optimization into the sampling process, enabling efficient and concise enforcement of hard constraints while maintaining high sample quality. Importantly, CCFM guarantees feasibility in the same manner as conventional repeated projection and, despite operating directly on noisy intermediate samples, is theoretically equivalent to projecting onto the feasible set defined by clean samples. Numerical experiments show that CCFM achieves superior performance in PDE solution generation and molecular docking.

Jinhao Liang
University of Virginia
jliang@email.virginia.edu

MS118

Generative PDE Solving from Partial Observation and Optimization-free PDE Constraints Enforcement

We present a unified framework for solving partial differential equations (PDEs) using video-inpainting diffusion transformer models. Unlike existing methods that devise specialized strategies for either forward or inverse problems under full or partial observation, our approach unifies these tasks under a single, flexible generative framework. Specifically, we recast PDE-solving as a generalized inpainting problem, e.g., treating forward prediction as inferring missing spatiotemporal information of future states from initial conditions. To this end, we design a transformer-based architecture that conditions on arbitrary patterns of known data to infer missing values across time and space. Our method proposes pixel-space video diffusion models for fine-grained, high-fidelity inpainting and conditioning, while enhancing computational efficiency through hierarchical modeling. Extensive experiments show that our video inpainting-based diffusion model offers an accurate and versatile solution across a wide range of PDEs and problem setups, outperforming state-of-the-art baselines. Moreover, we further develop a novel technology to ensure that the output solutions follow the underlying PDEs. We make an algorithmic improvement over the popular diffusion posterior sampling approach by adopting a learned gradient descent approach for robust inference time constraints enforcement.

Jeong Joon Park
University of Michigan
jjparkcv@umich.edu

MS118

Physics-Informed Posterior Sampling in Function

Spaces

Inverse problems in scientific computing are foundational yet challenging due to high dimensionality, ill-posedness, and uncertainty. Traditional methods like optimization or MCMC often struggle with computational expense and robustness to noise. While generative models have emerged for PDE solving, they often lack flexibility or data efficiency. In this talk, we explore physics-informed posterior sampling using flow-based models in function spaces, enabling resolution-agnostic, physics-consistent inverse solvers. We introduce Function-space Diffusion Posterior Sampling (FunDPS), the first infinite-dimensional diffusion framework for PDE inverses. By training unconditional denoisers on functions and guiding via sparse measurements with theoretical guarantees, FunDPS achieves up to 32% accuracy gains and 4x computational efficiency over grid-based baselines. The architecture also unlocks powerful multi-resolution strategies for further speedups. Building on this, we present the Decoupled Diffusion Inverse Solver (DDIS), which decouples prior learning from physics-informed sampling. This design mitigates oversmoothing in diffusion posterior sampling and enables the incorporation of advanced operator learning for inverse solving. We hope to provide new insights for infinite-dimensional generalization of posterior sampling, fostering scalable physics-informed tools for science.

Jiachen Yao
California Institute of Technology
jiachen.yao@caltech.edu

MS119

Decision Making and Control with Digital Twins for Underground Energy Storage

Decision making for underground energy storage must balance operational performance with safety under sparse observations and substantial subsurface uncertainty. We present a digital twin for uncertainty quantification (UQ) and closed-loop operations optimizatione.g., selecting injection schedules, pressure-management strategies, and monitoring actionsdriven by probabilistic state estimates from time-lapse geophysical data. The digital twin couples physics-based simulators with amortized Bayesian inference to learn a conditional posterior over reservoir states (e.g., pressure and saturation) given seismic monitoring data. These posterior distributions quantify uncertainty and propagate it into downstream control decisions. We then formulate risk-aware objectives that optimize operations subject to explicit safety criteria, using probability-of-failure (PoF) constraints and conditional value-at-risk (CVaR) penalties to manage low-probability, high-impact events such as reservoir overpressure that could compromise storage integrity.

Felix Herrmann
Georgia Institute of Technology
felix.herrmann@gatech.edu

Abinav Gahlot, Haoyun Li
Georgia Tech
agahlot8@gatech.edu, hli853@gatech.edu

MS119

Scientific Machine Learning Enhanced Forecasting in Subsurface Flows

Accurate forecasting of subsurface flow migration and trapping remains a central challenge in geoscience. Despite well-established physical understanding and advanced numerical simulators, predictive uncertainty persists due to geological heterogeneity, data scarcity, and the prohibitive cost of high-fidelity simulations. This talk presents a scientific machine learning framework to enhance forecasting capabilities for subsurface flows across scales. Starting from lab-scale experiments, we develop digital twins that combine experimental data, physics-based models, and machine learning surrogates to reproduce and predict multiphase flow dynamics under both pressure-driven and capillarybuoyancydominated flow regimes. The framework enables quantitative studies of energy storage retention and flow transport, while identifying key geological and barrier properties that control storage/operation performance. Through collaborations, this research aims to establish a scalable and interpretable SciML foundation for digital twins of subsurface energy storage, bridging experimental observables with predictive modeling of subsurface flow in realistic heterogeneous formations.

Hannah Lu
University of Texas at Austin
hannah.lu@austin.utexas.edu

MS120

Kernel Methods for Operator Learning

We present a kernel-based framework for learning operators between Banach spaces. We show that even with simple kernels, our approach is competitive in terms of cost-accuracy trade-off and either matches or beats the performance of Neural Network methods on a majority of PDE-based benchmarks. Additionally, our framework offers several advantages inherited from kernel methods: simplicity, convergence guarantees, a priori error estimates, and Bayesian UQ. We also present recent work on high-accuracy operator learning for the emulation of numerical solvers. We show that by emulating the Newton updates of a numerical solver, we can achieve machine precision on PDE forward and inverse problems using operator learning.

Matthieu Darcy
Caltech
mdarcy@caltech.edu

MS120

Data-driven Probabilistic Medium-Range Weather Forecasting

The recent revolution in data-driven methods for weather forecasting has lead to a fragmented landscape of architectures and training strategies, obscuring the drivers of forecast accuracy. Here, we demonstrate that state-of-the-art probabilistic skill requires neither architectural constraints nor specialized training heuristics. We introduce a scalable framework for learning atmospheric dynamics by combining a downsampled latent space with a history-conditioned local projector that resolves high-resolution physics. We find that our framework is robust to the choice of probabilistic estimator, supporting stochastic interpolants, diffusion models, and CRPS-based ensemble training. Validated against the Integrated Forecasting System and the deep learning probabilistic model GenCast, our framework achieves statistically significant improvements on most variables. These results suggest scaling a general-purpose model is sufficient for medium-range prediction, eliminating the need for tailored train-

ing recipes and proving effective across the spectrum of probabilistic frameworks. This is a joint work with Jean Kossaifi, Nikola Kovachki, Morteza Mardani, Daniel Leibovici, Suman Ravuri, Ira Shokar Edoardo Calvello, Mohammad Shoaib Abbas, Peter Harrington, Ashay Subramaniam, Noah Brenowitz, Boris Bonev, Wonmin Byeon, Karsten Kreis, Dale Durran, Arash Vahdat, Mike Pritchard and Jan Kautz.

Daniel Leibovici
NVIDIA
dleibovici@nvidia.com

MS120

Training Distribution Design in the Space of Probability Measures

What probability distribution should training data be sampled from to best approximate a target function or operator? This talk provides an answer in the setting that best refers to out-of-distribution (OOD) generalization accuracy with respect to a family of downstream tasks. The talk proposes to minimize average-case OOD accuracy, or its upper bound, over the space of probability measures. The minimizing data distribution depends not just on the model class (e.g., kernel regressors, neural networks) and the accuracy metric, but also on the target map itself. This leads to two implementable adaptive and target-dependent data selection algorithms based on either bilevel or alternating optimization. The new approach produces trained surrogate models that display increased robustness when evaluated on test data from shifted and unseen distributions. These models further empirically outperform others trained on traditional nonadaptive or target-independent data distributions in several function approximation, operator learning, and inverse problem tasks.

Nicholas H. Nelsen
Cornell University
nnelsen@cornell.edu

MS120

Discretization-Invariant Operator Learning: Algorithms and Theory

Learning operators between infinitely dimensional spaces is an important learning task arising in wide applications in machine learning, data science, mathematical modeling and simulations, etc. This talk introduces a new discretization-invariant operator learning approach based on data-driven kernels via deep learning. Compared to existing methods, our approach achieves attractive accuracy in solving forward and inverse problems, prediction problems, and signal processing problems with discretization invariant, i.e., networks trained with a fixed data structure can be applied to heterogeneous data structures without expensive re-training. Under mild conditions, quantitative generalization error will be provided to understand operator learning in the sense of non-parametric estimation.

Haizhao Yang
University of Maryland, College Park
haizhaoyang@yahoo.com

MS122

Integrating Prior Knowledge in Compositional Gaussian Processes

The fusion of physical knowledge with data-driven learning has allowed modelling practitioners to combine fundamental governing laws with the expressivity of machine learning, giving rise to hybrid models that have explanatory power where data is missing, whilst retaining the broad learning ability of machine learning technologies. There is thus great value in methods of this nature when attempting to model the behaviour of high value aerospace and civil assets, where data are often lacking, but one is still reliant on the flexibility of data-driven modelling. When constructing hybrid models, the form in which the domain insight exists is critical, as it governs the route by which knowledge is embedded into the model architecture. In many scenarios, this knowledge will not represent the entirety of the phenomena of interest, resulting in only a partial understanding of the driving physics. In this work, we consider how partial knowledge may be integrated into compositional, or deep Gaussian process models. In such model structures, individual Gaussian processes are stacked together as individual layers, where the output of one is fed as the input to another. Whilst embedding physical structure inside single layer Gaussian processes is a very active research area, less attention has been directed to integrating prior knowledge inside compositional architectures.

Matthew Jones
University of Sheffield
matthew.r.jones@sheffield.ac.uk

MS122

Uncertainty Quantification and Double Descent in Sparse Sensing

Sparse sensing is the paradigm of reconstructing full state of a complex system using few localized measurements and taking advantage of training data to propose a tailored basis specific to the system. The question of where to place sensors has been previously addressed with efficient techniques like QR pivoting and its extensions that account for engineering constraints. In this talk we turn to the question of the reconstruction algorithm and demonstrate that naive pseudoinverse leads to the phenomenon of double descent, or the instability of reconstruction the sensor and mode numbers match. We argue that optimal data-driven regularization of the reconstruction reduces the reconstruction variance, suppresses the double descent instability, and provides both an integral prediction of the error and a detailed spatial heatmap of how the error is distributed.

Andrei A. Klishin
University of Hawai?i at Manoa
aklishin@hawaii.edu

Andrei A. Klishin
University of Washington
aklishin@hawaii.edu

MS123

Hyper-Differential Sensitivity Analysis with Respect to Model Discrepancy

Computational models have made significant progress in simulating complex physical phenomena. However, this has led to a substantial increase in computational costs, making high-fidelity models impractical for analyses such as optimization. To address this challenge, low-fidelity models can be constructed, which are more amenable to optimization. Examples of low-fidelity modeling approaches include averaging small length scales, omitting physical processes, and employing data-driven model approxima-

tions. However, modeling errors propagate through the optimization process, resulting in suboptimal solutions. This work considers the use of hyper-differential sensitivity analysis to learn improved optimization solutions from limited high-fidelity model evaluations. We introduce a representation of model discrepancythe difference between high- and low-fidelity modelsand calibrate it in a Bayesian framework. The model discrepancy posterior is propagated through the optimization problem using post-optimality sensitivity analysis, yielding a posterior distribution for the optimization solution. We present an interpretable and efficient approach to generate optimization solution posterior samples. We demonstrate that as little as one high-fidelity solve can provide a significant improvement in the optimization solution. Furthermore, we leverage Bayesian optimal experimental design to guide high-fidelity model evaluations.

Joseph Hart
Sandia National Laboratories
joshart@sandia.gov

MS123

Detecting Model Form Error in Dynamical Systems Using Information Field Theory

Uncertainty in the governing equations is a major source of error when conducting inference in dynamical systems. Incorrect model forms cannot typically be compensated for by parameter calibration, and neglecting this issue risks biased estimates and misleading conclusions. While existing model selection techniques provide ways to choose among candidate models, they are often limited by a narrow pool of options that may not be representative of the unmodeled dynamics. An alternative line of work introduces correction terms into the governing equations, but most approaches rely on general forms, e.g., Gaussian processes or neural networks, that do not necessarily preserve physical constraints and symmetries. The goal of this presentation is to identify model correction terms from observations while ensuring a thermodynamically consistent parametrization. Our methodology builds on information field theory, a scalable Bayesian framework that enables us to embed the governing physics as prior knowledge over the field with a learnable trust. We then update the trust through a pretrained joint distribution of data and governing fields, using a diffusion-based model. We showcase the capability of our approach through synthetic examples spanning dynamical systems of different complexity, including Hamiltonian systems and those with energy dissipation. The results illustrate how the approach can recover missing physics while providing a principled quantification of model-form error.

Motahareh Mirfarah, Alex Alberts, Ilias Bilionis
Purdue University
mmirfara@purdue.edu, albert31@purdue.edu, ibilion@purdue.edu

MS123

Informing Predictive Uncertainties in the Presence of Model Misspecification: Generalized Bayesian Inference Approaches

In real-world engineering applications, computational models play a crucial role in guiding decision-making by providing predictions in regimes where observational data is unavailable. However, assessing the reliability of extrapolative predictions requires meaningful uncertainty quantification in calibration scenarios. While Bayesian inference provides a unifying framework for characterizing and extrapolating uncertainties, standard approaches can underpredict uncertainties when model misspecification is present. Even with the integration of physics-based or data-driven enhancements to reduce model discrepancy, the risk of misspecification remains. Advances in more generalized Bayesian inference techniques have demonstrated the potential to improve robustness to misspecification by employing more flexible loss function characterizations compared to conventional likelihood-based approaches. This work investigates novel formulations of such generalized losses that directly penalize poor uncertainty characterizations. We demonstrate this approach on an application problem in contaminant transport through heterogeneous porous media.

Rebekah White
Sandia National Laboratories
rebwhit@sandia.gov

MS124

Rl-Daunce: Reinforcement Learning-Driven Data Assimilation with Uncertainty-Aware Constrained Ensembles

Machine learning is increasingly used to enhance data assimilation. While supervised learning has become the standard approach, reinforcement learning (RL) with its sequential decision-making structure, aligns with the iterative nature of assimilation by dynamically balancing forecasts and observations. We present RL-DAUNCE (Reinforcement LearningDriven Data Assimilation with Uncertainty-Aware Constrained Ensembles) built on three core ideas: (i) RL-DAUNCE naturally inherits the computational efficiency of machine learning, while it uniquely structures its agents to mirror ensemble members in conventional data assimilation methods; (ii) multiple ensemble trajectories are advanced to quantify uncertainty beyond mean-state estimates; and (iii) constraints are enforced through a primaldual optimization strategy and action-space restrictions, ensuring physical consistency throughout the assimilation process. We apply RL-DAUNCE to the MaddenJulian Oscillation, a strongly intermittent atmospheric phenomenon with non-Gaussian behavior and multiple physical constraints. The method outperforms the standard ensemble Kalman filter, and matches the performance of constrained EnKF while requiring substantially less computation. RL-DAUNCE demonstrates strong skill in recovering intermittent signals, capturing extremes, and quantifying uncertainties, highlighting the potential of reinforcement learning as a flexible and efficient tool for constrained data assimilation.

Pouria Behnoudfar
University of Wisconsin Madison
p.behnoud@gmail.com

Nan Chen, Nan Chen
University of Wisconsin - Madison
chennan@math.wisc.edu, chennan@math.wisc.edu

MS124

MESS: A Multiproposal, Self-tuning, Dimension-free MCMC Algorithm with Applications to Bayesian PDE Inverse Problems

We present Multiproposal Elliptical Slice Sampling (MESS), a multiproposal extension of Elliptical Slice Sam-

pling (ESS) that allows multiple candidate states to be sampled in parallel at each iteration and favors accepting a candidate far from the current state. This can result in substantial reduction in wall clock time per independent sample, especially for expensive likelihoods. We also provide theoretical justification and experimental evidence for dimension-free mixing, showing that MESS is adequate for Bayesian PDE inverse problems. We demonstrate the algorithm on a toy model for solute transport. The model, strongly inspired by real physics, considers an implied infinite-dimensional parameter and a nonlinear parameter-to-observation map that generates highly non-Gaussian Bayesian posteriors with complex correlation structures.

Guillermina Senn
Norwegian University of Science and Technology
guillermina.senn@ntnu.no

MS125

Elastic Bayesian Model Calibration

Functional data are ubiquitous in scientific modeling. For instance, quantities of interest are modeled as functions of time, space, energy, density, etc. Uncertainty quantification methods for computer models with functional response have resulted in tools for emulation, sensitivity analysis, and calibration that are widely used. However, many of these tools do not perform well when the computer models parameters control both the amplitude variation of the functional output and its alignment (or phase variation). This paper introduces a framework for Bayesian model calibration when the model responses are misaligned functional data. The approach generates two types of data out of the misaligned functional responses: (1) aligned functions so that the amplitude variation is isolated and (2) warping functions that isolate the phase variation. These two types of data are created for the computer simulation data (both of which may be emulated) and the experimental data. The calibration approach uses both types so that it seeks to match both the amplitude and phase of the experimental data. The framework is careful to respect constraints that arise especially when modeling phase variation, and is framed in a way that it can be done with readily available calibration software.

Devin Francom
Los Alamos National Laboratory
dfrancom@lanl.gov

MS125

Scalable Non-Gaussian Variational Inference for Spatial Fields Using Sparse Autoregressive Normalizing Flows

We introduce a novel framework for scalable and flexible variational inference targeting the non-Gaussian posterior of a latent continuous function or field. For both the prior and variational family, we consider sparse autoregressive structures corresponding to nearest-neighbor directed acyclic graphs. Within the variational family, conditional distributions are modeled with highly flexible normalizing flows. We provide an algorithm for doubly stochastic variational optimization, achieving polylogarithmic time complexity per iteration. Empirical evaluations show that our method offers improved accuracy compared to existing techniques.

Paul Wiemann
The Ohio State University
wiemann.3@osu.edu

Matthias Katzfuss
University of WisconsinMadison
katzfuss@gmail.com

MS125

Spatial Warping Error Models for Addressing Structural Biases in Aerosol Climate Simulations

Aerosols influence the global climate by affecting cloud formation and altering the scattering and absorption of radiation. To reduce aerosol-related climate uncertainties, relevant climate model parameters are constrained by comparing the simulator outputs to observations of atmospheric conditions. Recent advances improve these comparisons by addressing the systematic model discrepancy between the simulator and atmospheric observations using a data-driven additive model discrepancy error. However, this treatment is not necessarily representative of typical structural errors in climate model outputs. In particular, misspecification in atmospheric dynamics can lead to spatial warping errors between the simulation and observations where the simulated structure of an atmospheric feature may be correct overall but it is displaced in space or distorted in shape. To address this, we propose a novel spatiotemporal modeling framework that models the spatial warping errors as random transport maps that capture these spatial distortions. We generate plausible transport maps using convex Gaussian processes and describe how simulation-based inference enables us to address the resulting likelihood-free inference problem for constrained spatial processes. When applied to the UKESM1 climate model and remotely sensed aerosol observations, our approach yields significantly higher likelihoods and stronger parameter constraints in comparison to models that rely solely on additive errors.

Erik Bensen, Victor Sanchez, Mikael Kuusela, Hamish Gordon
Carnegie Mellon University
ebensen@andrew.cmu.edu, vsanchez@andrew.cmu.edu, mkuusela@andrew.cmu.edu, hamishg@andrew.cmu.edu

MS125

Gaussian Processes for Uncertainty Quantification with Big Data

Uncertainty quantification is critical in machine learning for workflows such as error propagation, Bayesian optimization, and decision-making under uncertainty. While neural networks have predictive utility, they lack a simple and general approach to modeling uncertainty. In contrast, Gaussian processes (GPs) provide native uncertainties and allow explicit testing of model assumptions. However, standard GP methods face severe computational and memory bottlenecks with large datasets or high-dimensional inputs, and common isotropic or stationary covariance models are often too restrictive. We address the GP computational challenges in two stages. First, we introduce MuyGPs, a leave-one-out cross-validation framework that induces sparsity to enable scalable GP inference with a large number of observations. Second, we present a variable selection procedure and model designed for high-dimensional problems that mitigates the curse of dimensionality and can be combined with MuyGPs for computational efficiency. We demonstrate the effectiveness of this combined method in a few surrogate model cases, highlighting both their practical

advantages and their limitations.

Amanda Muyskens
Real Good AI
almuyskens@realgoodai.org

MS126

Scalable Langevin Monte Carlo Methods for Bayesian Inference

We introduce Langevin Monte Carlo methods for estimating expectations of observables under high-dimensional probability measures. We discuss discretization strategies and Metropolization methods for removing bias due to discretization error. We then present a new unbiased method for Bayesian posterior means based on kinetic Langevin dynamics that combines advanced splitting methods with enhanced gradient approximations. Our approach avoids Metropolis correction by coupling Markov chains at different discretization levels in a multilevel Monte Carlo approach. Theoretical analysis demonstrates that our proposed estimator is unbiased, attains finite variance, and satisfies a central limit theorem. We prove similar results using both approximate and stochastic gradients and show that our methods computational cost scales independently of the size of the dataset. Our numerical experiments demonstrate that our unbiased algorithm outperforms the gold-standard randomized Hamiltonian Monte Carlo. We then discuss whether it is always necessary to remove discretization bias in high-dimensional models and applications of unadjusted MCMC algorithms in Bayesian Neural Network models.

Peter Whalley
ETH Zurich
peter.whalley@math.ethz.ch

MS126

A Monte Carlo Framework for Calibrated Uncertainty Estimation in Sequence Prediction

Probabilistic prediction of sequences from images and other high-dimensional data remains a key challenge, particularly in safety-critical domains. In these settings, it is often desirable to quantify the uncertainty associated with the prediction (instead of just determining the most likely sequence, as in language modeling). In this paper, we consider a Monte Carlo framework to estimate probabilities and confidence intervals associated with sequences. The framework uses a Monte Carlo simulator, implemented as an autoregressively trained neural network, to sample sequences conditioned on an image input. We then use these samples to estimate probabilities and confidence intervals. Experiments on synthetic and real data show that the framework produces accurate discriminative predictions, but can suffer from miscalibration. To address this shortcoming, we propose a time-dependent regularization method, which produces calibrated predictions.

Qidong Yang
Massachusetts Institute of Technology
qidong@mit.edu

Weicheng Zhu, Laure Zanna, Tim Rudner, Carlos Fernandez-Granda
New York University
wz727@nyu.edu, lz1955@nyu.edu, tim.rudner@nyu.edu, cfgranda@cims.nyu.edu

MS126

Improved Energies and Wave Function Accuracy with Weighted Variational Monte Carlo

Neural network parametrizations have increasingly been used to represent the ground and excited states in variational Monte Carlo (VMC) with promising results. However, traditional VMC methods only optimize the wave function in regions of peak probability. The wave function is uncontrolled in the tails of the probability distribution, which can limit the accuracy of the trained wavefunction approximation. To improve the approximation accuracy in the probability tails, this paper interprets VMC as a gradient flow in the space of wave functions, followed by a projection step. From this perspective, arbitrary probability distributions can be used in the projection step, allowing the user to prioritize accuracy in different regions of state space. Motivated by this theoretical perspective, the paper tests a new weighted VMC method on the antiferromagnetic Heisenberg model for a periodic spin chain. Compared to traditional VMC, weighted VMC reduces the error in the ground state energy by a factor of 2 and it reduces the errors in the local energies away from the mode by large factors of 10^2–10^4.

Huan Zhang
École normale supérieure PSL
huan.zhang@lmd.ipsl.fr

MS127

Robust Optimal Experimental Design for Inverse Problems: Theory and Software Implementation

In this talk, I present a family of probabilistic methods for solving challenging model-based optimal experimental design (OED) problems. This includes robust and constrained binary OED problems with black-box objectives for batch sensor placement, and trajectory optimization for sequential decision making in large-scale inverse problems. In the probabilistic approach, the optimization variable is viewed as a random variable and is associated with a parametric conditional probability model. A stochastic optimization problem is then formulated and efficiently solved by following a policy gradient approach. I will also give a high-level overview of PyOED, a highly extensible scientific package that enables rapid development, testing, and benchmarking of model-based OED methods for inverse problems. The presented probabilistic approaches to OED are implemented and are publicly available through PyOED.

Ahmed Attia
Argonne National Laboratory
Mathematics and Computer Science Division
aattia@anl.gov

MS127

Power Flow Informed Priors for Energy Market Applications

Power flow models are central to energy market operations, forming the foundation of software tools ranging from simple line flow calculations to security-constrained optimal power flow. These tools underpin how Independent System Operators determine prices and how market participants engage in the energy markets. Yet the inherent uncertainty

in inputsload forecasts, renewable generation, and system topologyposes major challenges for both operational reliability and trading strategies. In this work, we propose constructing prior distributions over power flow quantities of interest, grounded in the empirical marginal distributions of relevant inputs pushed through power flow simulations. We study the resulting priors, with a particular focus on how their structure depends on both the data and the software implementation of the underlying physical models. This perspective highlights how software-informed priors can guide scientific machine learning in energy systems, enabling more robust decision-making under uncertainty.

Michael Brennan
Salvo Energy
mbrennan@soleaenergy.com

Matthew Parno
Solea Energy
parnomd@gmail.com

MS127

Active Learning Strategies for Modeling Unobserved Variables

Consider a complex physical system and two models that measure a Quantity of Interest(QoI) from the system. They produce similar results with different levels of accuracy and computational time, but with discrepancies in their inputs. Several surrogate models in Uncertainty Quantification and machine learning applications present this setup, such as: First, one model can depend on deterministic inputs, but the other on stochastic inputs. Next, the inputs can have different resolutions of the QoI, or the inputs share certain parameters but not the full span of input parameters in the system. Our goal is to model the discrepancy between models using a common set of input parameters and learn a model that decreases the sample complexity. There are three challenges: the number of unknown variables and the distribution are unknown. Sampling from the model with high accuracy is expensive. Next, given a set of known outputs, there are not uniquely associated. To this end, we learn a non-linear map from model one to model two. Constructed by a neural network using known and hidden parameters. We minimize a suitable loss function, defined by the mean-squared error and the Maximum Mean Discrepancy (MMD), which we later use to create an active learning strategy that finds more suitable samples to minimize it, subsequently learning the underlying distribution with fewer samples. Our numerical results show that our method outperforms state-of-the-art methods in modeling this task.

Juan Cardenas
University of Colorado Boulder
juca7891@colorado.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

Tony Zahtila
Stanford University
tzahtila@stanford.edu

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

MS127

Multi-Fidelity Approaches to Bayesian Optimal Experimental Design

Experimental data is often vital but expensive to collect, making it important to find experimental design conditions that maximize the value of these experiments. Bayesian optimal experimental design (OED) provides a framework for identifying the ideal design by leveraging a mathematical model that simulates the experimental outcomes. The practicality of OED is often limited by the *fidelity* of the model used. High-fidelity models are computationally expensive, making OED intractable, while low-fidelity models can lead to suboptimal designs. Furthermore, current OED methods often overlook the computational cost of the model evaluations themselves. This work introduces two multi-fidelity approaches to overcome these limitations. The first approach uses approximate control variate estimators to expedite the estimation of the expected information gain (EIG), providing EIG estimates that are unbiased with respect to the high-fidelity model, but with significant computational speedups. The second approach treats the multi-fidelity model ensemble as a set of numerical experiments that can also be planned via OED. This method uses a sequential and goal-oriented strategy to jointly learn the quantity of interest as well as the relationships between numerical and physical experiments in a cost-effective way. We demonstrate the value of these multi-fidelity approaches on a set of practical engineering applications benchmark problems.

Thomas Coons, Caleb Dahlke, Aniket Jivani
University of Michigan
tcoons@umich.edu, cdahlke@umich.edu,
ajivani@umich.edu

Tommie Catanach, Ruben Villarreal
Sandia National Laboratories
tacatan@sandia.gov, rubvill@sandia.gov

Xun Huan
University of Michigan
xhuan@umich.edu

MS128

Data Assimilation with Normalizing Flow Applied to Seismic Images

Seismic measurements governed by the wave equation provide an indirect means of estimating subsurface properties such as density and seismic velocity. In practice, uncertainty in active sources and environmental noise, together with the strong nonlinearity of wave propagation, render the associated inverse problems ill-posed. Moreover, inversion approaches often depend heavily on the choice of initial subsurface models, limiting their reliability for monitoring applications. We focus on the problem of monitoring injected subsurface CO2 to verify containment for long-term geological storage. The migration of CO2 in the subsurface is described by multiphase flow physics and depends on heterogeneous permeability and porosity fields. We apply two ensemble-based data assimilation approaches, comparing the ensemble Kalman filter (EnKF) to conditional normalizing flow filters (CNFFs). We first demonstrate a failure for each method on a toy Gaussian problem and the large seismic-CO2 system before directly comparing their

accuracy and uncertainty estimates on the large system.

Grant Bruer
Georgia Institute of Technology
gbruer@gatech.edu

Abhinav Gahlot
Georgia Tech
agahlot8@gatech.edu

Felix Herrmann
Georgia Institute of Technology
felix.herrmann@gatech.edu

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS128

Weakly-Constrained 4D Var for Downscaling with Uncertainty using Data-Driven Surrogate Models

Dynamic downscaling typically involves using numerical weather prediction (NWP) solvers to refine coarse data to higher spatial resolutions. Data-driven models such as FourCastNet have emerged as a promising alternative to the traditional NWP models for forecasting. Once these models are trained, they are capable of delivering forecasts in a few seconds, thousands of times faster compared to classical NWP models. However, as the lead times, and, therefore, their forecast window, increase, these models show instability in that they tend to diverge from reality. In this work, we propose to use data assimilation approaches to stabilize them when used for downscaling tasks. Data assimilation uses information from three different sources, namely an imperfect computational model based on partial differential equations (PDE), from noisy observations, and from an uncertainty-reflecting prior. In this work, when carrying out dynamic downscaling, we replace the computationally expensive PDE-based NWP models with FourCastNet in a "weak-constrained 4DVar framework" that accounts for the implied model errors. We demonstrate the efficacy of this approach for a hurricane-tracking problem; moreover, the 4DVar framework naturally allows the expression and quantification of uncertainty. We demonstrate, using ERA5 data, that our approach performs better than the ensemble Kalman filter (EnKF) and the unstabilized FourCastNet model, both in terms of forecast accuracy and forecast uncertainty.

Philip Dinenis
Argonne National Laboratory, USA
pdinenis@anl.gov

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS128

Identifying Informative Designs for Constraining Viscosity in MgO under Shocked Conditions

When modeling materials with extremely high strain rates, such as shock compression experiments, or materials with extremely low strain rates, such as rheological flows, solids are often modeled as a fluid. For example, Earths mantle, though mostly solid, is treated as a viscous fluid over geological timescales with resistance to flow characterized by an effective viscosity (Colvin et al., 2003; Mikaelian, 2013). However, experimental constraints on mantle-relevant materials, such as MgO, at appropriate pressures and temperatures remain limited (Foley, 2019; Foley et al., 2012; Karato, 2011). Prior work has explored using Richtmyer-Meshkov instability (RMI) growth rates to infer the strength of MgO under shock, comparing experimental results to theory to help constrain viscosity (Perez et al., 2025). We extend this by introducing optimal experimental design (OED) aimed at identifying experimental conditions that best constrain MgO viscosity under shock loading. Our setup considers a shock incident on a corrugated interface between epoxy and MgO, with design variables including shock strength, initial amplitude, and wavelength. We define the experimental objective in terms of maximizing expected information gain on viscosity.

Codie Kawaguchi, Xun Huan, Eric Johnsen
University of Michigan
cfkawa@umich.edu, xhuan@umich.edu,
ejohnsen@umich.edu

Sonya Dick
South Dakota School of Mines and Technology
sonya.dick@sdsmt.edu

MS129

Deep Gaussian Process Poisson Modeling of Large-Scale Satellite Count Data for Computer Model Calibration

The National Aeronautics and Space Administration (NASA) launched the Interstellar Boundary Explorer (IBEX) satellite in 2008 to learn more about the region at the boundary between our solar system and interstellar space, called the heliopause. IBEX detects energetic neutral atoms (ENAs) originating from the heliopause and enables space scientists to estimate their rate of emission throughout the sky. Space scientists have developed a variety of competing theories (with corresponding computer simulations) to explain ENA generation and propagation and hope to use data collected by IBEX to evaluate the relative strength of evidence for each proposed model. Statistical computer model calibration allows data from IBEX to be paired with computer simulations to accomplish this task. However, the count observations, limited but data-intense simulation runs and field measurements, and nonstationary response surface present some unique challenges. We propose a Vecchia-approximated deep Gaussian process Poisson model for the rate of ENAs. Our model has shown a greater ability to learn the complex, nonstationary mean response surface, and provide appropriate uncertainty quantification in different regions of the input space, while maintaining a simpler covariance function. We integrate this model into a novel Markov chain Monte Carlo computer model calibration framework to understand the distribution of the parameters that govern ENA generation.

Steven D. Barnett
Virginia Polytechnic Institute & State University
sdbarnett@vt.edu

Robert Gramacy
Virginia Tech
rbg@vt.edu

Lauren Beesley, Dave Osthus
Los Alamos National Laboratory
lvandervort@lanl.gov, dosthus@lanl.gov

Yifan Huang
University of Alabama in Huntsville
yh0017@uah.edu

Fan Guo, Daniel Reisenfeld
Los Alamos National Laboratory
guofan@lanl.gov, dreisenfeld@lanl.gov

MS129

Functional Data Analysis and Deep Learning for Accelerated Statistical Model Calibration

We present a unified framework for the robust calibration of complex physical and engineering models under uncertainty. Our methodology integrates functional data analysis, deep learning surrogates, and optimization-based uncertainty calibration in function space. Data are represented via wavelet expansions, preserving the temporal structure of high-dimensional data throughout the surrogate construction and calibration processes. A Monte Carlo dropout neural network is trained to map model input parameters to functional response coefficients, providing accurate surrogate predictions with epistemic uncertainty bands. Calibration is performed through deterministic optimization of model parameters against data profiles, enabling parameter estimation informed by entire functional curves rather than isolated pointwise measurements. By coupling concepts from functional data analysis and deep learning, our framework establishes a scalable foundation for future wavelet-enabled statistical calibration.

Kwesi A. Ohene-Obeng, Luis Damiano, Khachik Sargsyan
Sandia National Laboratories
kaohene@sandia.gov, ladamia@sandia.gov,
ksargsy@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS129

Neural Bayesian Methods for Spatio-Temporal Data Fusion of Remote Sensing Observations

Hyperspectral and multispectral imaging satellites provide complementary remote sensing data across diverse temporal, spatial, and spectral resolutions. However, these resolutions are often misaligned: satellites with high temporal resolution typically have low spatial resolution, and vice versa. Statistical fusion of multi-sensor data offers the potential to generate high-resolution images across space, time, and spectra at near-global scales. Here, we present a highly scalable Neural Bayesian framework for fusing massive hyperspectral datasets from instruments with varying temporal, spatial, and spectral support. Our approach leverages a Gaussian process formulation in which 3-D Convolutional Neural Networks with Residual Blocks estimate posterior distributions of kernel parameters, while sparse matrix techniques enable efficient approximate Bayesian inference. By combining deep learning, sparse methods modeling, and high-performance computing, our method delivers well-calibrated uncertainties with orders-of-magnitude speedups over existing approaches, enabling rapid, high-fidelity spatio-temporal data fusion at remarkable scales.

Joaquim Teixeira
University of California, Los Angeles
joaquimteixeira95@gmail.com

Amy Braverman
Jet Propulsion Laboratory
California Institute of Technology
Amy.Braverman@jpl.nasa.gov

Maggie Johnson
Jet Propulsion Laboratory
maggie.johnson@jpl.nasa.gov

Sudipto Banerjee
Department of Biostatistics
UCLA
sudipto@ucla.edu

MS130

Improved Subseasonal Forecasting of Extreme Weather Using Statistical Machine Learning

Extreme temperatures cause more than $100 billion in damages annually within the United States alone. This means it is of great societal importance to improve predictive skill of these temperature extremes on subseasonal timescales (2-6 weeks) beyond that of the traditional physics-based model. In this study, we evaluate the performance of statistical-machine learning (ML) methods, including Echo State Networks (ESN) and Random Forests (RF), for forecasting weekly average temperatures across five U.S. regions, with forecast horizons ranging from one to four weeks. In addition to point forecasts, we assess each model's ability to predict extremes, defined as weekly temperature anomalies exceeding one standard deviation from climatology. All ML models are benchmarked against two standard baselines: a persistence model and a climatology-plus-trend model. We additionally compare feature importance between models. Our results provide a comprehensive comparison of ML methods for subseasonal temperature forecasting and offer new insights into the complexity of feature relevance across lead times and regions. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525. SAND2025-09300A.

Katherine Goode, Maike Holthuijzen
Sandia National Laboratories
kjgoode@sandia.gov, mfholth@sandia.gov

Jacob Johnson
Sandia National Laboratories
University of Wisconsin-Madison
jacjohn@sandia.gov

Meredith G. L. Brown
Sandia National Laboratories
merbrow@sandia.gov

Thomas Ehrmann
SANDIA National Laboratories
tsehrma@sandia.gov

MS130

Locally Adaptive Conformal Inference for Operator Models

Operator models are regression algorithms between Banach spaces of functions. They have become an increasingly critical tool for spatiotemporal forecasting and physics emulation, especially in high-stakes scenarios where robust, calibrated uncertainty quantification is required. We introduce Local Sliced Conformal Inference (LSCI), a distribution-free framework for generating function-valued, locally adaptive prediction sets for operator models. We prove finite-sample validity and derive a data-dependent upper bound on the coverage gap under local exchangeability. On synthetic Gaussian-process tasks and real applications (air quality monitoring, energy demand forecasting, and weather prediction), LSCI yields tighter sets with stronger adaptivity compared to conformal baselines. We also empirically demonstrate robustness against biased predictions and certain out-of-distribution noise regimes.

Trevor Harris
University of Connecticut
zft24002@uconn.edu

MS130

Bayesian Closure Modeling for Dynamic Systems

Computational simulations of dynamical systems often require constitutive or closure models to represent unresolved phenomena, enhance computational efficiency, or correct model form error in the simulation. Data-driven closure models (DDCMs) employ statistical machine learning (ML) to learn the closure term and uncertainty as a function of dynamic system states and have the potential to increase the predictive capability of computational simulations. Our approach addresses learning dynamics in data sparsity regimes by augmenting the observed measurements with dense state and closure terms estimated from experimental data using a particle filter. For trustworthiness, accurate uncertainty quantification (UQ) for DDCM predictions is essential. Uncertainties that arise from sparse or noisy experimental data are mapped to the estimated closure term using the particle filter, and then probabilistic statistical models propagate uncertainties to the predicted quantity. We demonstrate our methodology on an exemplar from dynamic ecological modeling and provide preliminary results using Bayesian hierarchical models with regularization for the DDCM. SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Toryn Schafer, Dan Drennan
Texas A&M University
tschafer@tamu.edu, danjdrennan@tamu.edu

Kyle Neal, Rileigh Bandy, Moe Khalil, Teresa Portone
Sandia National Laboratories
kneal@sandia.gov, rjbandy@sandia.gov, mkhalil@sandia.gov, tporton@sandia.gov

MS130

An Overview of Hybrid Statistical and Deep Neural Models for Spatial and Spatio-Temporal Processes

Environmental processes include complex dependencies across time, space, and sub-process. These processes are inherently dynamic, where spatial processes evolve with time. It has been a long-standing problem in spatio-temporal statistics to develop efficient, flexible, and realistic data-driven models for such processes in a way that appropriately accounts for uncertainty in data, process, and parameters. With the popularization of deep neural models in machine learning, a large collection of new methods and computational approaches are available for statisticians and applied mathematicians to use to address these challenges. In this overview lecture, I will present some hybrid deep neural/statistical approaches that have been used in environmental statistics in recent years. These include methods in which statistical models are embedded within deep learning paradigms and methods in which deep neural models are embedded in statistical models.

Christopher Wikle
Department of Statistics
University of Missouri
wiklec@missouri.edu

MS131

Numerical Analysis of Discontinuous Galerkin Method for Nonlinear Stochastic Convection-Diffusion Equation

There has been a growing interest for designing efficient high order numerical algorithms to approximate solutions of stochastic partial differential equations (SPDEs). In this talk, we develop fully-discrete high-moment stability analysis and error estimates for nonlinear stochastic convection-diffusion equations in one and higher dimensions, adopting local discontinuous Galerkin (LDG) method in space and implicit-explicit Euler-Maruyama scheme in time. Existing work on DG methods for SPDEs mostly focus on linear equations and semi-discrete analysis. This is because solutions of SPDEs are generally not time differentiable and are not bounded in the path variable. In this work, we prove high-moment stability analysis and (path-wise) error estimates for fully nonlinear equation on subsets whose probability converge to one. Numerical experiments are provided to verify the strong order of accuracy of our numerical schemes.

Yiming Chen
Ohio State University
chen.11042@osu.edu

MS131

Diffusion-Based Stochastic Operator Networks for Uncertainty Quantification in Stochastic Partial Differential Equations

We propose a novel framework for effectively quantifying the uncertainty in the solutions of stochastic partial differential equations (SPDEs) using stochastic operator learning. SPDEs has been used intensively to model complicated physical systems under uncertainty, however, while accurate, classical solvers require problem-specific discretizations, adaptive meshing near singular features, which are often tedious and costly. To avoid such difficulties, we develop the Stochastic Operator Network (SON) by combining the stochastic optimal control mechanism of the Stochastic Neural Network (SNN) with the architecture of the DeepONet to learn the noisy solution operators

of SPDEs directly from the available data. By formulating the training procedure as a stochastic optimal control problem, the SON is able to learn the uncertainty presented in the solution operators through the diffusion coefficients appearing in the SNN. Numerical experiments on benchmark SPDEs demonstrate the accuracy and robustness of the proposed method in replicating several noisy solution operators.

Phuoc Toan Huynh
Florida State University
th24ba@fsu.edu

MS131

Uncertainty Quantification and Prediction of Two-Dimensional Geophysical Turbulence with Generative Models

Two-dimensional geophysical turbulence exhibits strong intermittency, non-Gaussian variability, and finite-memory effects that challenge deterministic reduced-order closures and yield over-confident forecasts. This talk presents a generative-AI framework for probabilistic prediction and uncertainty quantification (UQ) in 2D geophysical flows (e.g., quasi-geostrophic turbulence). We combine conditional diffusion models with physics-aware embeddings to produce calibrated ensemble forecasts directly in vorticity/streamfunction space. The models are trained to approximate the conditional law of subgrid tendencies given resolved states (and short histories), with soft constraints enforcing incompressibility and energy/enstrophy budgets. We quantify both aleatoric and epistemic uncertainty via likelihood/score objectives and Bayesian weight regularization and validate with a-posteriori rollouts against high-resolution simulations. Diagnostics include spectra and fluxes, PDFs of vorticity and gradients, extreme-event exceedance probabilities, and long-horizon stability.

Xuping Xie
Old Dominion University
xxie@odu.edu

MS131

Learning Gradient Flow Dynamics from Particle Data

We propose an interpretable learning framework for identifying the gradient flow dynamics of particle-based evolution systems. The objective is to recover the underlying energy functional, which naturally decomposes into internal, drift, and interaction terms. The key idea of our approach is to design a learning algorithm that preserves the structure of the energy. The inferred model can then be used to generate more particle data over longer time horizons and under various initial conditions. This approach yields a physically grounded model that not only captures observed dynamics but also generalizes to predict behaviors in complex particle systems.

Qihao Ye
University of California, San Diego
q8ye@ucsd.edu

MS132

Microstructural Models for Coarsening Dynamics in Polycrystalline Materials

Polycrystalline materials consist of small grains separated by boundaries with different lattice orientations, and these structures strongly affect material properties. A method by which the grain structure can be engineered in polycrystalline materials is through grain growth (coarsening) of a starting structure. The evolution of the grain/grain boundary structure and associated grain growth is a very complex, multiscale, and multiphysics process. It involves the dynamics of grain boundaries, triple junctions in 2D (triple curves/lines and quadruple points in 3D), and the dynamics of lattice misorientations. So, grain growth can be regarded as the anisotropic evolution of a large cellular network and can be described by a set of deterministic local evolution laws for the growth of individual grains combined with stochastic models for the interaction between them. In this talk, I will present recent perspectives on mathematical modeling, simulation and analysis of the evolution of the grain boundary network in polycrystalline microstructures. Relevant experiments and future research will also be discussed. This work is collaborated with Katayun Barmak, Batuhan Bayir, William M Feldman, David Kinderlehrer, Chang (Kamala) Liu, Chun Liu, Masashi Mizuno, Thuong Nguyen, Jeffrey Rickman, and Matthew Patrick, and is supported by the DMREF program(DMS-2118172) and the Simons Foundation Fellowship Award SFI-MPS-SFM-00010667.

Yekaterina Epshteyn
Department of mathematics
University of Utah
epshteyn@math.utah.edu

MS132

A Model-free Method for Discovering Symmetry in Differential Equations

Symmetry in differential equations reveals invariances and offers a powerful means to reduce model complexity. Lie group analysis characterizes these symmetries through infinitesimal generators, which provide a local, linear criterion for invariance. However, identifying Lie symmetries directly from scattered data, without explicit knowledge of the governing equations, remains a significant challenge. In this talk, I will introduce a numerical scheme that approximates infinitesimal generators from data sampled on an unknown smooth manifold, enabling the recovery of continuous symmetries without requiring the analytical form of the differential equations. We employ a manifold learning technique, Generalized Moving Least Squares, to prolongate the data, from which a linear system is constructed whose null space encodes the infinitesimal generators representing the symmetries. Convergence bounds for the proposed approach are derived. Several numerical experiments, including ordinary and partial differential equations, demonstrate the methods accuracy, robustness, and convergence, highlighting its potential for data-driven discovery of symmetries in dynamical systems.

Max Kreider
Penn State University
mbk6295@psu.edu

MS132

PDE Foundation Model: Generalization, Learning to Learn and More

This talk explores the path toward foundation models for scientific computing, focusing on learning solution operators for parametric partial differential equations (PDEs). The core of the presentation showcases our contributions to overcoming these limitations: (1) A multi-modal trans-

former framework (PROSE) that achieves zero-shot generalization to new PDE families by jointly processing input functions and symbolic equations; and (2) A meta-learning strategy (LeMON) that enables a single pre-trained neural operator to rapidly adapt to new physical systems with minimal data. Together, these advances bridge the gap between specialized, single-task solvers and a flexible, multi-operator foundation model capable of real-time simulation and design.

Jingmin Sun
Johns Hopkins University
jsun143@jhu.edu

MS133

Steering Generative Models with Experimental Data for Protein Fitness Optimization

Designing proteins with desired quantitative properties requires searching an astronomically large sequence space while collecting only limited experimental data. Recent protein generative modelssuch as diffusion and language modelsoffer promising ways to guide this search. Yet most existing methods rely on surrogate rewards or large labeled datasets, leaving open how well they perform when only a few hundred wet-lab measurements are available. In this talk, I will present a systematic study of protein fitness optimization in this data-scarce regime and introduce Steered Generation for Protein Optimization (SGPO), a framework that couples generative protein models with small amounts of experimental data to efficiently discover high-fitness protein variants. SGPO encompasses multiple types of generative models and steering mechanisms, enabling us to identify the most effective design choices for practical protein engineering campaigns.

Raul Astudillo
Mohamed bin Zayed University of Artificial Intelligence
raul.astudillo@mbzuai.ac.ae

MS133

Scaling Bayesian Optimization of Composite Functions for High-Dimensional Inverse Problems: Application to Electron Microscopy

We study high-dimensional inverse problems of the form:

$$\min_{x \in \mathcal{X}} g(h(x); y),$$

where $h(x)$ is a d-dimensional output of an expensive derivative-free simulation under parameter x, y is the ground truth vector, and g is a known function that measures discrepancy between simulated and the true outputs. The objective is to efficiently find parameter values x that minimize this error. Bayesian Optimization of Composite Functions (BOCF) tackles such problems by modeling $h(x)$ with a multi-output Gaussian process (GP) surrogate and leveraging the known structure g, improving query efficiency. However, when the dimensionality d grows large, inference for multi-output GPs becomes computationally expensive and makes BOCF infeasible. In this work, we propose a scalable BOCF framework that aggregates components of $h(x)$ to create lower-dimensional outputs, reducing computational cost while preserving the advantage in query efficiency offered by the composite structure. We demonstrate the approach on electron microscopy, where $h(x)$ represents simulated images under parameters x (specimen thickness and misorientation) and y denotes experimental images. Numerical results show that our method recovers the parameters x more accurately and with fewer simulations than standard Bayesian optimization and at a fraction of the runtime required by the original BOCF.

Dasol Yoon
Department of Materials Science and Engineering
Cornell University
dy327@cornell.edu

Poompol Buathong
Cornell University
pb482@cornell.edu

Chia-Hao Lee
School of Applied and Engineering Physics
Cornell University
cl2696@cornell.edu

David A. Muller
School of Applied and Engineering Physics, Cornell Universit
Kavli Institute at Cornell for Nanoscale Science, Cornell Un
david.a.muller@cornell.edu

Peter I. Frazier
School of Operations Research and Information Engineering
Cornell University
pf98@cornell.edu

MS133

IAGO: Implicit Adaptive Generative Design Optimization, a Data-Scarce Optimization Framework for Physics-Based, Automatically Manufacturable Designs

We present IAGO, a framework for design optimization. The method uses latent conditional neural fields trained to represent shapes as signed distance functions (SDFs) that that satisfy manufacturability and geometry constraints. After training, Bayesian Optimization (BO) in the latent space identifies high- and low-performing designs by calling a physics solver to compute domain-specific metrics. Thereafter, a refined model is trained with the additional objectives of generating solutions that resemble high-performing designs, while discouraging recurrence of low-performing ones. This two-steps cycle of model training and BO is repeated until performance requirements are met. Building on prior work "[Berzins et al., Geometry-Informed Neural Networks]", IAGO integrates knowledge of prior designs and is efficient in data-scarce regimes. Further, in contrast to other solver-in-the-loop methods ("[Radler et al., Diverse Topology Optimization using Modulated Neural Fields]", "[Hoyer et al., Neural reparameterization improves structural optimization]"), IAGO does not rely on gradients of the cost function, as common in topology optimization. Our experimental results in the context of structural mechanics show that IAGO produces high-performing and manufacturable designs after a few cycles. In this talk, we will provide a few examples of IAGO's application to structural mechanics problems.

Orazio Pinti
Atomic Machines
orazio.pinti@atomicmachines.com

Leonardo Ferreira Guilhoto
University of Pennsylvania
guilhoto@sas.upenn.edu

Simone Venturi
Atomic Machines
simone@atomicmachines.com

Paris Perdikaris
University of Pennsylvania
pgp@seas.upenn.edu

Marta D'Elia
Sandia National Laboratories
marti.delia@gmail.com

MS133

Probabilistic Learning of the Purkinje Network from the 12 Lead Electrocardiogram

The identification of the Purkinje conduction system in the heart is a challenging task, yet essential for a correct definition of cardiac digital twins for precision cardiology. Here, we propose a probabilistic approach for identifying the Purkinje network from non-invasive clinical data such as the standard electrocardiogram (ECG). We use cardiac imaging to build an anatomically accurate model of the ventricles; we algorithmically generate a rule-based Purkinje network tailored to the anatomy; we simulate physiological electrocardiograms with a fast model; we identify the geometrical and electrical parameters of the Purkinje-ECG model with Bayesian optimization and approximate Bayesian computation. The proposed approach is inherently probabilistic and generates a population of plausible Purkinje networks, all fitting the ECG within a given tolerance. In this way, we can estimate the uncertainty of the parameters, thus providing reliable predictions. We test our methodology in physiological and pathological scenarios, showing that we are able to accurately recover the ECG with our model. We propagate the uncertainty in the Purkinje network parameters in a simulation of conduction system pacing therapy. Our methodology is a step forward in creation of digital twins from non-invasive data in precision medicine.

Francisco Sahli Costabal
Pontificia Universidad Catolica de Chile
fsahlic@uc.cl

MS134

Multi-Stage Latent Space Dynamics Identification

Numerical solutions of partial differential equations (PDEs) are often computationally expensive, driving the need for efficient reduced-order models (ROMs). Gaussian Process-based Latent Space Dynamics Identification (GPLaSDI) is a recently proposed data-driven, non-intrusive ROM framework that compresses simulation data using an autoencoder and learns a system of user-chosen ordinary differential equations (ODEs) governing the latent space dynamics. This allows for rapid predictions by interpolating and evolving the low-dimensional ODEs in the latent space. Additionally, using GPs to interpolate ODEs in the latent space allows for efficient uncertainty-based adaptive sampling of the parameter space. While GPLaSDI has produced effective ROMs for numerous problems, the autoencoder may have difficulty accurately reconstructing training data while also satisfying the imposed dynamics in the latent space, particularly in complex or high-frequency regimes. To address this, we propose multi-stage Latent Space Dynamics Identification (mLaSDI), which trains GPLaSDI models in sequential stages. Each stage learns a new decoder to correct the residual error from previous stages, improving prediction accuracy and reducing sensitivity to hyperparameter choices. We demonstrate mLaSDI on a multiscale oscillatory problem, 2D vortex shedding, and the 1D-1V Vlasov equation, achieving high prediction accuracy for all problems. LLNL-CONF-2006046

William Anderson
Lawrence Livermore National Laboratory
anderson316@llnl.gov

MS134

Learning Dynamical Systems from Time and Frequency-Responsedata

Dynamical systems are a principal tool in the modeling, prediction, and control of physical phenomena with applications ranging from structural health monitoring to electrical power network dynamics, from heat dissipation in complex microelectronic devices to vibration suppression in large wind turbines. Direct numerical simulation of these mathematical models may be the only possibility for accurate prediction or control of such complex phenomena. However, in many instances, a high-fidelity mathematical model describing the dynamics is not readily available. Instead, one has access to an abundant amount of input/output data via either experimental measurements or a black-box simulation. The goal of data-driven modeling is, then, to accurately model the underlying dynamics using input/output data only. In this talk, we will investigate various approaches to data driven modeling of dynamical systems using systems-theoretical concepts. We will consider both frequency-domain and time-domain measurements of a dynamical system including parametrically varying dynamics. In some instances, we will have true experimental data, and in others we will have access to simulation data. We will illustrate these concepts in various examples ranging from structural dynamics to microelectromechanical systems.

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS134

Understanding The Effectiveness of Lossy Compression in Machine Learning and AI Training Sets

Often datasets for AI used in scientific applications can be large extreemly large making them difficult to move and store. One approach to dealing with this are data reduction techniques including sampling (that reduces the number of elements of data) and dimensionality reduction (that reduces the number of features that each element has) but also lossy data compression which reduces the footprint of data by reducing precision of each feature. In this work, we explore how to most effectively use modern error bounded lossy compressors to reduce the footprint of these datasets to improve the storage and transfer of such datasets.

Robert R. Underwood
Clemson University
School of Computing

runderwood@anl.gov

Jon Calhoun
Clemson University
Holcombe Department of Electrical and Computing Engineering
jonccal@clemson.edu

Sheng Di, Franck Cappello
Argonne National Laboratory
sdi1@anl.gov, cappello@anl.gov

MS135

Transport Vs. MCMC for Conditional Sampling

A common task in uncertainty quantification is modeling a conditional distribution for parameters of interest given data. To compute descriptive statistics, one often requires independent samples from this distribution. The standard tools are Markov Chain Monte Carlo (MCMC) methods, which provide strong theoretical guarantees but often suffer from slow convergence and require re-initialization for each new conditioning variable. Measure transport methods, by contrast, construct deterministic or stochastic maps between reference and target measures, offering amortized sampling and scalability. In this talk, we present a systematic comparison of measure transport and MCMC across a suite of benchmark UQ problems. Specifically, we construct triangular maps via flow-based and stationary transport models to demonstrate the efficacy of such a conditional sampling approach. Through our experiments, we highlight measure transport as a practical and powerful alternative to MCMC.

Alex Johnson-Vázquez, Bamdad Hosseini
University of Washington
verano13@uw.edu, bamdadh@uw.edu

Ricardo Baptista
University of Toronto
r.baptista@utoronto.ca

MS135

Epistemic Uncertainty Quantification with Applications in Machine Learning

Epistemic uncertainty arises from limited knowledge about the underlying data-generating process and plays a central role in the reliability, robustness, and safety of machine learning systems. In this talk, I will discuss recent advances in quantifying epistemic uncertainty, with a focus on approaches grounded in imprecise probability and related frameworks. I will highlight how these methods go beyond conventional probabilistic modelling, capturing ambiguity and ignorance in a principled way. Applications in machine learning including out-of-distribution generalisation, compositional learning, and AI safety will illustrate the potential of epistemic uncertainty quantification to enable more trustworthy and rational decision-making.

Krikamol Muandet
CISPA - Helmholtz Center for Information Security
muandet@cispa.de

MS135

Principal Curves in Metric Spaces And The Space of Probability Measures

We introduce principal curves in Wasserstein space, and in general compact metric spaces. Our motivation for the Wasserstein case comes from optimal-transport-based trajectory inference, where a developing population of cells traces out a curve in Wasserstein space. Our framework enables new experimental procedures for collecting high-density time-courses of developing populations of cells: time-points can be processed in parallel (making it easier to collect more time-points). However, then the time of collection is unknown, and must be recovered by solving a seriation problem (or one-dimensional manifold learning problem). We propose an estimator based on Wasserstein principal curves, and prove it is consistent for recovering a curve of probability measures in Wasserstein space from empirical samples. This consistency theorem is obtained via a series of results regarding principal curves in compact metric spaces. In particular, we establish the validity of certain numerical discretization schemes for principal curves, which is a new result even in the Euclidean setting.

Andrew Warren, Anton Afanassiev, Forest Kobayashi, Young-Heon Kim
University of British Columbia
awarren@math.ubc.ca, anton.a@math.ubc.ca, fkobayashi@math.ubc.ca, yhkim@math.ubc.ca

Geoffrey Schiebinger
The University of British Columbia
geoff@math.ubc.ca

MS135

Proximal Transport Divergences for Robust Generative Flows and Bayesian Inference

Proximal optimal transport (OT) divergences are a novel class of discrepancy measures that interpolate between information divergences and OT distances through infimal convolution. These divergences provide a principled foundation for the proximal optimization methods commonly used in generative modeling. We establish that these divergences possess key mathematical properties: smoothness, boundedness, and computational tractability. We then demonstrate that generative flows trained with these divergences are well-posed with desirable regularity properties, making them particularly effective for learning distributions on low-dimensional manifolds while maintaining stable and robust training dynamics. Finally, we apply these robust generative models to likelihood-free inference for generating samples from Bayesian posteriors supported on low-dimensional manifolds.

Benjamin Zhang
University of North Carolina at Chapel Hill
bjz@unc.edu

MS136

Data Driven Modeling Methods to Enable Real-Time Digital Twins

We present a numerical framework for constructing a targeted digital twin (tDT) that directly models the dynamics of quantities of interest (QoIs) in a full digital twin. The proposed approach employs memory-based flow map learning (FML) to develop a data-driven model of the QoIs. The trained FML model provides a compact and efficient surrogate representation of the unknown system, enabling the

real-time prediction and control for the full DT. We demonstrate the effectiveness of our approach through examples in computational fluid dynamics.

Qifan Chen
Department of Mathematics
The Ohio State University, Columbus
chen.11010@osu.edu

MS137

Learning Pseudo-Differential Operators

Pseudo-differential operators (PDOs) are a broad and fundamental class of operators in analysis and partial differential equations, with wide-ranging applications in mathematical and computational sciences. In this talk, I will present our recent work on learning elliptic PDOs from noisy data. Within a wavelet–Galerkin framework, we formulate this task as a structured infinite-dimensional regression problem with multiscale sparsity. Building on this structure, we propose a sparse, data- and computation-efficient estimator that leverages a novel matrix compression scheme tailored to the learning task and a nested-support strategy to balance approximation and estimation errors. In addition to establishing convergence rates for the estimator, we show that the learned operator induces an efficient and stable Galerkin solver whose numerical error matches its statistical accuracy. Our results therefore contribute to bringing together operator learning, data-driven solvers, and wavelet methods in scientific computing.

Jiaheng Chen
University Of Chicago
jiaheng@uchicago.edu

Daniel Sanz-Alonso
University of Chicago
sanzalonso@uchicago.edu

MS137

Scaling Scientific Machine Learning: Integrating Theory and Numerics in Both Training and Inference

Scaling scientific machine learning (SciML) requires overcoming bottlenecks at both training and inference. On the training side, we study the statistical convergence rate and limits of deep learning for solving elliptic PDEs from random samples. While our theory predicts optimal polynomial convergence for PINNs, optimization becomes prohibitively ill-conditioned as networks widen. By adapting descent strategies to the optimization geometry, we obtain scale-invariant training dynamics that translate polynomial convergence into concrete compute and yield compute-optimal configurations. On the inference side, I will introduce Simulation-Calibrated SciML (SCaSML), a physics-informed post-processing framework that improves surrogate models without retraining or fine-tuning. By enforcing physical laws, SCaSML delivers trustworthy corrections (via Feymann-Kac simulation) with approximate confidence intervals, achieves faster and near-optimal convergence rates, and supports online updates for digital twins. Together, these results integrate theory and numerics to enable predictable, reliable scaling of SciML in both training and inference.

Yiping Lu
Courant Institute, New York University
Northwestern University
yiping.lu@northwestern.edu

MS137

On the Benefits of Active Data Collection in Operator Learning

We study active data collection strategies for operator learning when the target operator is linear and the input functions are drawn from a mean-zero stochastic process with continuous covariance kernels. With an active data collection strategy, we establish an error convergence rate in terms of the decay rate of the eigenvalues of the covariance kernel. We can achieve arbitrarily fast error convergence rates with sufficiently rapid eigenvalue decay of the covariance kernels. This contrasts with the passive (i.i.d.) data collection strategies, where the convergence rate is never faster than linear decay ($\sim n^{-1}$). In fact, for our setting, we show a *non-vanishing* lower bound for any passive data collection strategy, regardless of the eigenvalues decay rate of the covariance kernel. Overall, our results show the benefit of active data collection strategies in operator learning over their passive counterparts.

Unique Subedi
University of Michigan
subedi@umich.edu

MS137

Theory-to-Practice Gap for Operator Learning

The theory-to-practice gap describes the discrepancy between the parametric complexity and the sampling complexity of learning. For a set of functions with a fixed parametric complexity, or the best possible approximation error rate with respect to the number of model parameters, we examine the corresponding sampling complexity, or the best error rate with respect to the number of data samples given a reconstruction algorithm. In some cases, such as polynomial or trigonometric polynomial approximation, these two rates coincide: determining N degrees of freedom requires N point samples. However, for approximation spaces of neural networks, this intuition does not hold; indeed, the sampling complexity rate is fundamentally limited. As the parametric complexity rate becomes arbitrarily fast, the optimal sampling complexity rate stays uniformly bounded. In this work, we extend the theory-to-practice gap result for learning in finite dimensions to the infinite-dimensional setting of operator learning.

Margaret K. Trautner, Samuel Lanthaler
California Institute of Technology
mtrautne@caltech.edu, slanth@caltech.edu

Phillip Grohs
University of Vienna
philipp.grohs@univie.ac.at;

MS138

The RTO Map for Hierarchical Sampling

In many hierarchical inverse problems, not only do we want to estimate high- or infinite-dimensional model parameters in the parameter-to-observable maps, but we also have to estimate hyperparameters that represent critical assumptions in the statistical and mathematical modelling processes. As a joint effect of high-dimensionality, nonlinear dependence, and nonconcave structures in the joint posterior distribution over model parameters and hyperparame-

ters, solving inverse problems in the hierarchical Bayesian setting poses a significant computational challenge. In this talk, we will present scalable optimization-based Markov chain Monte Carlo (MCMC) methods for solving hierarchical Bayesian inverse problems with nonlinear parameter-to-observable maps and a broader class of hyperparameters. Our algorithmic development is based on the recently developed scalable randomize-then-optimize (RTO) method for exploring the high- or infinite-dimensional parameter space. By using RTO either as a proposal distribution in a Metropolis-within-Gibbs update or as a biasing distribution in the pseudo marginal MCMC, we present efficient sampling tools for hierarchical Bayesian inversion. In particular, the integration of RTO and the pseudo marginal MCMC has sampling performance robust to model parameter dimensions. With minor modifications, the presented methods can also be applied to problems with the Poisson likelihood and other non-Gaussian observation processes.

Tiangang Cui
Monash University
tiangang.cui@sydney.edu.au

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

MS138

Excursions in Bayesian Treatments of Model Error

Berliner et al. (2023) discusses some challenges and strategies for treating model error in a Bayesian context. They distinguish between two classes of models. The first targets physically-based statistical models that are sufficiently manageable to allow incorporation of a stochastic "model error process'. The second addresses large-scale models in which incorporation of a model error process and updating its distribution is impractical. Current work extends some of the ideas brought forth in Berliner et al. (2023). Under the hierarchical Bayesian modeling framework, we discuss avenues for learning and quantifying the uncertainty in the dimension-reduction operator and for addressing the process model / model error confounding issues. We illustrate the ideas using an application inspired by a classical oceanographic inverse problem.

Radu Herbei
The Ohio State University
herbei@stat.osu.edu

MS139

Uncertainty-Aware Sensor Optimization for Safety-Critical Systems

Sensor placement within safety-critical systems such as nuclear reactors and aerospace applications is highly constrained by limited capacity, extreme operating conditions, high costs, restricted accessibility, and stringent safety regulations. Real-time monitoring of temperature, pressure, velocity, and power requires accurate field reconstruction from sparse measurements while quantifying noise-induced uncertainties. This work develops an uncertainty-aware framework for optimizing constrained sensor locations and field reconstruction through reduced-order models, D-optimal design criteria, followed by QR pivoting. Reconstruction results from QR-optimized sensors demonstrate high accuracy, reduced noise-induced uncertainties, and reconstruction error reduction with minimal sensor requirements across both experimental and simulation datasets. While this approach achieves high reconstruction accuracy for linear systems, we further develop a mutual information-based sensor selection pipeline to address nonlinear reconstruction challenges. This approach demonstrated accurate field reconstruction by selecting sensors that maximize mutual information with the full-field, applied to a boundary shock mixing layer configuration.

Niharika Karnik
Idaho National Lab
niharika.karnik@inl.gov

MS139

Machine LearningBased Uncertainty Quantification for Condition Monitoring in Nuclear Power Plant Generation Risk Analysis

This work presents a machine learningbased uncertainty quantification framework for condition monitoring and generation risk analysis in nuclear power plants. The framework integrates machine learning models that estimate equipment reliability from condition monitoring datawhile also explicitly accounting for associated uncertaintiesin combination with techniques for synthetically generating data on electricity prices and plant output such that the statistical characteristics (i.e., data uncertainty) of the original data are preserved. These uncertainty sources are then propagated through a dynamic nuclear power plant generation risk model to quantify their combined impact on generation risk. The proposed approach improves predictions of future lost generation and supports cost-effective, risk-informed decision-making for project prioritization and long-term equipment reliability planning.

Congjian Wang
Idaho National Lab
congjian.wang@inl.gov

MS140

Data Assimilation and Discrepancy Modeling with Shallow Recurrent Decoders

Complex physical systems are often modeled through simulations to approximate their dominant dynamics. However, these simulation models could inevitably neglect small-scale or hidden processes or oversimplify parameter correlations, leading to reconstructions that can diverge from reality. This creates a pressing need for data assimilation, the process of integrating simulation data with predictive models to produce coherent and accurate estimates of the full state of complex physical systems. In this paper, we propose a framework for Data Assimilation with SHallow REcurrent Decoder (DA-SHRED) which bridges the sim-to-real gap between simulations and sensor data for real physics. For real-world systems on high-dimensional spatiotemporal fields, where the full state cannot be directly observed and must be inferred from sparse sensor measurements, we leverage the latent space learned from a reduced simulation model via SHRED networks, and update these latent variables using real sensor data to accurately reconstruct the full system state. Furthermore, our algorithm incorporates a SINDy-based regression model in the latent space to identify functionals corresponding to missing dynamics in the reduced simulation model. In presented examples, DA-SHRED successfully recovers missing dynamics in highly volatile systems, demonstrating that the combination of efficient temporal encoding and physics-informed correction enables robust data assimilation under

sparse sensing constraints.

Yuxuan Bao
Univeristy of Washington
baoyx@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

MS140

Unpacking Model Inadequacy: Applications to Nuclear Fuels and Materials

We develop a Bayesian calibration framework for nuclear fuels and materials applications that explicitly represents model–data mismatch via a Gaussian-process inadequacy. The approach propagates uncertainty from parameters, experimental noise, and model form, with a hierarchical variance structure to capture group and cross-group noise. For a nuclear graphite property prediction application, where we are interested in predicting the irradiation-induced dimension change, creep, Young's modulus change ratio, and coefficient of thermal expansion change ratio, we obtain average predictive-error reductions of 54%, 65%, 17%, and 17% when inadequacy is included. We illustrate engineering impact with a multiphysics model of a very-high-temperature reactor prismatic reflector brick, analyzing stresses under high fluence and temperature. Accounting for model inadequacy markedly improves predictive accuracy and provides a robust basis for reliable graphite component design in advanced reactors. For a tristructural isotropic (TRISO) particle application, we compare traditional Arrhenius equation fitted to experimental data and the more advanced lower-length-scale (LLS)-informed model, which considers microstructure information. Applying the inverse UQ process on the experimental datasets revealed modeling inadequacy to be the most dominant source of uncertainty. Experimental noise uncertainty is also significant; however, model parameter uncertainty can be considered negligible. Interestingly, both the Arrhenius equation and the LLS-informed model demonstrated similar levels of modeling inadequacy. For the forward predictive UQ on test datasets, the KOH framework improved both the accuracy and quality of quantified uncertainties in comparison to the standard Bayesian framework. This is true for both the Arrhenius equation and the LLS-informed model. In comparing these modeling approaches, both demonstrated similar performance at the engineering scale, while the LLS-informed model expectedly outperformed the Arrhenius equation at the mesoscale. Overall, across both the applications, it is important to explicitly account for modeling inadequacy in the UQ process, and reinforce the need for continuous refinement of physics-based models in order to address the modeling inadequacy.

Som L. Dhulipala
Idaho National Laboratory
Som.Dhulipala@inl.gov

Parikshit Bajpai, Pierre-Clément Simon, Jacob Hirschhorn
Idaho National Lab
parikshit.bajpai@inl.gov, pierreclement.simon@inl.gov, jacob.hirschhorn@inl.gov

Benjamin Spencer
Idaho National Laboratory
benjamin.spencer@inl.gov

Stephen Novascone
Idaho National Lab
stephen.novascone@inl.gov

MS140

Model-Form Uncertainty Quantification and Management in Hierarchical Systems

Many scientific and engineering problems involve predictions that depend on uncertain informationfrom noisy data to incomplete physical models. Characterizing, propagating, and predicting the performance of complex systems with hierarchical structure is particularly challenging in the presence of uncertainties and modeling errors. In this talk, I will present the Extended Polynomial Chaos Expansion (EPCE) framework, a robust and efficient paradigm that simultaneously propagate all sources of uncertainties and errors in order to provide a probabilistic representation of specific quantities of interest (QoI). We characterize the full probability distribution of the QoI and the uncertainty in the failure probability pertaining to its tails. By combining EPCE with kernel density estimation (KDE) and directional derivatives, we construct sensitivity measures that connect these statistical metrics of QoI to the model-form error and data error to assess their individual and combined impacts. We demonstrate the presented approach in several engineering systems with hierarchical interactions.

Zhiheng Wang
Texas Tech University
Civil, Environmental, and Construction Engineering
zhiheng.wang@ttu.edu

MS141

A Probabilistic Token Attribution Score For Large Language Models

The generative nature of current Large Language Models (LLMs) is expressed by the conditional probabilities that LLMs compute to sample each response token conditioned on the previous tokens. These probabilities encode the distributional structure that the model learns in training and exploits in inference. We have designed a model-agnostic probabilistic token attribution measure using Bayes rule to invert the logits so as to capture the model's internal representation of the distribution over token sequences. The representation is independent of details of model computational structure. This representation yields the conditional probability of the response given the prompt, and of the response given the prompt with a token marginalized away. Our attribution score is the log of the ratio of these probabilities. We further compute entropy of token distributions, conditioned on the entire remaining context. The interplay between entropy and attribution score sheds light on LLM behavior. We exhibit the behavior of 7 models on 6 prompts and identify token sensitivity and response stability, thereby improving interpretability and guiding users to focus on uncertain or unstable parts of the generation.

Carlo Graziani, Shilpika, Bethany Lusch
Argonne National Laboratory
cgraziani@anl.gov, shilpika@anl.gov, blusch@anl.gov

Michael E. Papka
Argonne National Laboratory
University of Chicago

papka@anl.gov

Venkatram Vishwanath
Argonne Leadership Computing Facility, Argonne National Lab.
venkatv@mcs.anl.gov

MS141

Quantifying Uncertainty in a Poorly Specified Methane Emissions Inverse Problem

Methane, the main component of natural gas, is the second-largest contributor to climate change after carbon dioxide. Methane has a higher heat-trapping potential but shorter lifetime than carbon dioxide, and therefore, rapid reduction of methane emission can have quick and large climate change mitigation impacts. Reducing emissions from oil and gas production facilities, turns out to be a particular promising avenue due to technological advances in continuous emission monitoring technology. We present a Bayesian framework for quick emission detection and localization using continuous methane concentration data measured by multiple monitoring sensors on oil and gas production facilities, demonstrate its effectiveness under real-world conditions and discuss ideas for comprehensive uncertainty quantification that incorporates source localization uncertainty.

Dorit Hammerling
Colorado School of Mines
hammerling@mines.edu

William Daniels
John Hopkins University
wdanie16@jh.edu

MS141

Generative Modeling for Bayesian Inverse Problems

Operational data processing pipelines in applications like atmospheric remote sensing are based on repeatedly solving Bayesian inverse problems involving computationally expensive forward models. In addition to computational costs, the resulting posterior distribution is usually non-Gaussian even when using normal priors, due to moderately non-linear forward models. The reported uncertainty is still based on a Laplace approximation of the posterior covariance, which is known to be prone to underestimating uncertainties. In this work, we apply generative modeling, specifically diffusion based models, to learn the joint distribution of model inputs and outputs, and use this technique to solve the underlying inverse problem fast while providing comprehensive description of the posterior, thanks to the conditional inference provided by inverting the diffusion flow.

Otto Lamminpaeae
Jet Propulsion Laboratory
otto.m.lamminpaa@jpl.nasa.gov

MS142

Improving Scalability of Multilevel Markov Chain Monte Carlo Using Multiple Proposals

Model-based simulations are an important tool for predicting physical phenomena. In the presence of uncertain physical parameters, which describe the system of interest, accurately assessing confidence in these results, e.g., performing Bayesian inference, can be computationally expensive, particularly when the uncertain physical parameters are spatially varying quantities. Not only are the corresponding high-fidelity (fine grid) simulations expensive, but commonly used statistical approaches, such as Markov chain Monte Carlo (MCMC), require an exceedingly high number of simulations. In this work, we develop a new scalable sampling method which complements the algorithmic framework of multilevel MCMC, where coarse grid simulations are used to inform the fine level proposal distribution, thereby accelerating MCMC. Specifically, using tools from algebraic multigrid, we form a (scalable) fine grid Gaussian random field realization from a fine grid proposal by combining Gaussian random fields sampled across multiple levels of discretization. To address the deterioration of coarse level scaling (which occurs in such settings), we leverage a combination of level-dependent redistribution and MCMC proposals performed in parallel. In this talk, we describe this new approach, corresponding theory, and numerical results when applied to a 3D subsurface flow application.

Hillary Fairbanks
Lawrence Livermore National Laboratory
fairbanks5@llnl.gov

Panayot Vassilevski
Portland State University
panayot@pdx.edu

MS142

Statistical Emulation with Input Dimension Reduction for Complex Forward Models

Gaussian process (GP) modeling in high-dimensional input space can result in inefficient emulators and is computationally challenging for a given budget. To tackle these issues, we introduce a fully Bayesian framework that seamlessly integrates dimensionality reduction with GP modeling and inference. Our approach, built on a hierarchical Bayesian model with priors on the Stiefel manifold. Additionally, we introduce a Deep Gaussian Processes (DGP) with built-in dimension reduction, providing a more flexible and powerful tool for complex datasets. We demonstrate that the proposed method improves predictive performance and uncertainty quantification, providing a principled and robust alternative to existing methods.

Alex Konomi, Eric Herrison Gyamfi
University of Cincinnati
alex.konomi@uc.edu, gyamfien@mail.uc.edu

Emily L. Kang
Department of Mathematical Sciences
University of Cincinnati
kangel@ucmail.uc.edu

MS142

Addressing Parameter Non-Identifiability Using Bayesian Calibration on the Active Subspace

The inference of model parameters for moderate- to high-dimensional model systems can be challenging, since the full set of parameters is often unidentifiable in the sense that they are not uniquely determined by the available data, particularly for applications with restrictive data constraints. Additionally, the computational cost of performing Bayesian inference in a high-dimensional space may

not allow for adequate exploration of the entire parameter domain. These issues may be mitigated by using active subspace techniques to identify the subspace of identifiable inputs, then exploiting this subspace to explore only those directions in the input space that will actively inform parameter estimates. Furthermore, such an approach does not require one to determine which subsets of parameters will be fixed versus estimated prior to the inference procedure. In this talk, we illustrate how the use of such a subspace-based approach enables us to inform the maximum number of identifiable parameters (often a larger number than the number of dimensions utilized for inference), while still maintaining convergence properties of MCMC chains. Additionally, we demonstrate the use of this methodology on a model of tumor growth dynamics for which the original parameter set is unidentifiable.

Allison L. Lewis
Lafayette College
lewisall@lafayette.edu

Rebecca Everett
Department of Mathematics and Statistics
Haverford College
reverett@haverford.edu

MS142

All Emulators are Wrong, Many are Useful, and Some are More Useful Than Others

Accurate and efficient surrogate modeling is essential for modern computational science, and there are a staggering number of emulation methods to choose from. With new methods being developed all the time, comparing the relative strengths and weaknesses of different methods remains a challenge due to inconsistent benchmarking practices and (sometimes) limited reproducibility and transparency. In this work, we present a large-scale, fully reproducible comparison of 29 distinct emulators across 60 canonical test functions and 40 real emulation datasets. To facilitate rigorous, apples-to-apples comparisons, we introduce the R package `duqling`, which streamlines reproducible simulation studies using a consistent, simple syntax, and automatic internal scaling of inputs. This framework allows researchers to compare emulators in a unified environment and makes it possible to replicate or extend previous studies with minimal effort, even across different publications. Our results provide detailed empirical insight into the strengths and weaknesses of state-of-the-art emulators and offer guidance for both method developers and practitioners selecting a surrogate for new data. We discuss best practices for emulator comparison and highlight how `duqling` can accelerate research in emulator design and application.

Kelin Rumsey, Devin Francom
Los Alamos National Laboratory
knrumsey@lanl.gov, dfrancom@lanl.gov

Graham Gibson
LANL
gcgibson@lanl.gov

MS143

Variational Bayesian Inference in UncertaintyQuantification.jl

Bayesian updating plays a pivotal role in engineering disciplines such as structural health monitoring, geotechnical engineering, and system identification to update probabilistic models based on collected data. The posterior distribution describing the updated knowledge on the uncertainties is typically analytically intractable and characterised only through sampling methods based on Markov Chain Monte Carlo (MCMC), which can be computationally prohibitive. Stochastic transport maps for Bayesian inference offer a compelling alternative, overcoming this limitation by reformulating the Bayesian updating problem as a deterministic, invertible mapping between a simple reference distribution and the target posterior. This variational problem is then solved through sparse integration techniques to efficiently approximate the full posterior probability density function with significantly reduced likelihood evaluations compared to traditional MCMC. This work presents the latest developments on variational Bayesian inference in the freely available Julia package UncertaintyQuantification.jl. Based on the strong foundation of the framework the implemented algorithms enable the parameter updating of complex high fidelity models. We present applications to two engineering problems, one involving interfacing with a third party solver, and compare the accuracy and computational demand to classic MCMC methods.

Lukas Fritsch
Institute for Risk and Reliability
Leibniz University Hannover
fritsch@irz.uni-hannover.de

Jasper Behrensdorf
Institute for Risk and Reliability
Leibniz Universität Hannover
behrensdorf@irz.uni-hannover.de

Ander Gray
UK Atomic Energy Authority (UKAEA)
ander.gray@ukaea.uk

Matteo Broggi
Institute for Risk and Reliability
Leibniz University Hannover
broggi@irz.uni-hannover.de

Michael Beer
Leibniz University Hannover
beer@irz.uni-hannover.de

MS143

UQLab - Recent Updates and a Rich Upcoming Roadmap

UQLab is a general purpose uncertainty quantification framework that has recently turned 10 since its first public beta release in July 2015. With multiple state-of-the-art modules ranging from surrogate modeling to advanced statistical inference, from reliability to sensitivity analysis, it is an ever evolving, actively developed software tool that has been adopted by thousands of users worldwide. In this talk, we want to showcase some of the recent and upcoming additions to the software package, which include:

- A powerful surrogate modeling benchmarking infrastructure, that allows researchers to compare their algorithms against the state-of-the-art on different case studies in a fair and standardized way;
- A wide selection of new surrogate models, including multiple machine-learning methods such as neural networks and foundation models;

- The ability of handling time-series data and models, including advanced surrogate models for dynamic data;
- A connection to UM-Bridge, a powerful interface to highly increase interoperability between computational models and uncertainty quantification tools.

Stefano Marelli
Chair of Risk, Safety and Uncertainty Quantification
Institute of Structural Engineering, ETH Zurich
marelli@ibk.baug.ethz.ch

Adéla Hlobilová
Chair of Risk, Safety and Uncertainty Quantification
ETH Zurich
adela.hlobilova@ibk.baug.ethz.ch

Bruno Sudret
ETH Zurich
Chair of Risk, Safety and Uncertainty Quantification
sudret@ibk.baug.ethz.ch

MS143

PyTUQ: UQ Tools and Workflows for Computational Models

The Python Toolkit for Uncertainty Quantification (PyTUQ, https://github.com/sandialabs/pytuq) is a lightweight Python library for a range of uncertainty quantification tasks and workflows. Features include conventional tools such as polynomial chaos machinery with mixed bases, global sensitivity analysis, quadrature point generation, linear regression, Bayesian inference with various flavors of Markov chain Monte Carlo. PyTUQ also includes advanced methods such as Bayesian compressed sensing, sampling-based Rosenblatt transformation and embedded model error calibration. In this work, we will present the structure of the library, provide several educational examples, and demonstrate PyTUQ's application to various scientific problems, ranging from Earth system modeling to chemistry and materials science.

Khachik Sargsyan, Emilie Grace Ed Baillo
Sandia National Laboratories
ksargsy@sandia.gov, eebaill@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

MS144

Multi-Fidelity Budget Constrained Optimal Experimental Design

In many scientific and engineering domains, data collection can be expensive and time-consuming, particularly when high-fidelity simulations or experiments are involved. Multifidelity modeling offers a powerful solution by combining data from multiple sources of varying accuracy and cost, potentially enabling more efficient learning. This work introduces a multi-fidelity Gaussian Process framework to jointly model many fidelities of data for optimal experimental design under strict budget constraints. The focus is on developing dynamic, cost-aware policies that guide the sequential selection of experiments to maximize an information-based utility. A central challenge addressed in this work is how to effectively represent the state of the design process in settings where decision sequences and resource usage evolve dynamically. These contributions are broadly applicable and provide a foundation for future extensions to more complex, constrained decision-making settings.

Caleb Dahlke, Thomas Coons, Xun Huan
University of Michigan
cdahlke@umich.edu, tcoons@umich.edu,
xhuan@umich.edu

Tommie Catanach
Sandia National Laboratories
tacatan@sandia.gov

MS144

A Bilinear Inverse Problem with Forward Operator Inaccuracy Applied to Optical Imaging

Linear inverse problems are common in real-world applications from industry to medical imaging, but handling inaccuracies in the associated forward operators is a relatively unstudied problem. In this work, we assume a set of candidate forward operators and suggest principal component analysis for modeling their variation from the mean, which leads to a bilinear tensor inverse problem. We present two optimization algorithms and Gibbs sampling for approximately solving the problem. As a real-world test case, we apply the algorithms to account for the anatomical inaccuracy that is present in the sensitivity profiles or Jacobian matrices in diffuse optical tomography when an atlas-based model of the head geometry is used instead of the subject's own anatomical model. We demonstrate improvements in the spatial localization and contrast-to-noise-ratio in example reconstructions of simulated brain activity. In addition, Bayesian optimal experimental design that is robust to anatomical variations is discussed for optode placement in diffuse optical tomography.

Aada Hakula, Pauliina Hirvi, Nuutti Hyvonen
Aalto University
aada.hakula@aalto.fi, pauliina.hirvi@aalto.fi, nuutti.hyvonen@aalto.fi

MS144

Model-Misspecification-Aware Bayesian Optimal Experimental Design

Bayesian Optimal Experimental Design (BOED) seeks to select an experimental design that provides the most informative data for estimating unknown parameters in the context of Bayesian inverse problems. All frameworks of BOED, however, rely on the assumption of a well-specified model i.e. the model used for inference is capable of capturing the true distributions of the data and its relationship with the parameters of interest. In real-world applications this is however rarely the case: errors or approximations are a natural part of the process. For BOED specifically, this means optimizing (or over-optimizing) experiment design based on wrong assumptions. In this paper we discuss a setting in which several linear Gaussian models are available to represent a certain phenomenon of interest and it is not possible to assess with certainty which is the best one to use. For inversion, a single model choice will have to be made (due to computational cost, for example). However, for the optimal sensor choice, we propose a methodology that can take this uncertainty among the input models into account and lead to a "misspecification-aware" sensor choice. We show how this sensor choice produces, on average, posterior distributions with less bias and variance

than those obtained with a classic, well-specified, BOED formulation.

Andrea Scarinci
Massachusetts Institute of Technology
scarinci@mit.edu

Ricardo Babtista
University of Toronto
Canada
r.baptista@utoronto.ca

Han Cheng Lie
University of Potsdam
Institute of Mathematics
hanlie@uni-potsdam.de

MS144

Nonlinear Oed with Column Subset Selection

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

Amit Subrahmanya
Argonne National Lab
famitnagesh@anl.gov

MS145

Control-Oriented Optimal Sensor Placement for Infinite-Dimensional Bayesian Inverse Problems

We propose a control-oriented optimal experimental design (cOED) framework for infinite-dimensional Bayesian linear inverse problems. The focus is on designing sensor placements in the service of optimal control problems where uncertain parameters must be inferred from sensor data. Conventional Bayesian OED methods typically generate sensor placements that minimize the uncertainty in the inferred parameter. However, such designs are often suboptimal for demands of the optimal control problem. In contrast, our cOED approach seeks experimental designs that minimize the uncertainty in either the controlled state variable or the control objective. We establish a mathematical framework for uncertainty quantification and cOED applicable to parameterized PDE-constrained control problems with linear dependencies on the control and inversion parameters. Furthermore, we develop scalable computational methods for determining these control-oriented sensor placements and for assessing the resulting uncertainty in the control objective. We demonstrate the effectiveness of the proposed approach in the context of a model problem motivated by heat transfer applications.

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

Madhusudan Madhavan, Arvind Saibaba
North Carolina State University
mmadhav@ncsu.edu, asaibab@ncsu.edu

Bart G. Van Bloemen Waanders, Rebekah White
Sandia National Laboratories
bartv@sandia.gov, rebwhit@sandia.gov

MS145

An Adaptive Ensemble Smoother Multiple Data Assimilation (ES-MDA) Method for Carbon Capture Storage Reservoirs

We present an adaptive Ensemble Smoother Multiple Data Assimilation (ES-MDA) strategy within the context of carbon capture and storage (CCS) reservoirs, where the goal is characterization of subsurface field properties. In this study, we compare our findings to state-of-the-art approaches for iterating the ES to convergence. Our results include inversion performance metrics of the proposed adaptive scheme and existing strategies on a 2D field-scale reservoir injection setup. In geologic carbon sequestration applications, observations can include field measurements such as well pressure history as well as derived products from seismic acquisitions and cross-well electrical resistivity measurements. A variety of these data can then be combined to infer reservoir features such as the evolution of CO 2 gas-plume spread through time. Informed risk assessment is critical to the success of proposed reservoir development projects. Ensemble techniques address this issue in providing access to the approximate posterior probability distribution function (pdf), enabling sampling of matched history models, and evaluation of uncertainty. Our primary contribution is an adaptive formulation of ES-MDA that reduces the number of required data assimilation steps and corresponding forward model evaluations, while still obtaining a meaningful UQ.

Kyle Ivey
Institute of Geophysics and Planetary Physics,
Scripps Institute of Oceanography
krivey@ucsd.edu

Matthias Morzfeld
Scripps Institution of Oceanograpghy
University of California San Diego
matti@ucsd.edu

Chaoyi Wang, Joshua A. White, Chris Sherman,
Christina Morency
Lawrence Livermore National Laboratory
wang140@llnl.gov, white230@llnl.gov,
sherman27@llnl.gov, morency1@llnl.gov

Robert Mellors
Scripps Institution of Oceanograpghy
University of California San Diego
rmellors@ucsd.edu

MS145

Solving the Stochastic Inverse Problem As Optimization Over the Probability Measure Space

We investigate the stochastic inverse problem: Given ρ_y, how to find ρ_x such that $\rho_y = G_{\#}\rho_x$, where probability distributions ρ_x and ρ_y are connected via the forward map G. Depending on G's property, the solution can have drastically different behavior. In the overdetermined case, we formulate a variational problem $\min_{\rho_x} D(G_{\#}\rho_x, \rho_y)$, and find that different choices of the metric D significantly affect the quality of the reconstruction. When D is set to be the Wasserstein distance, the reconstruction is the marginal distribution, while setting D to be a ϕ-divergence reconstructs the conditional distribution. In the underde-

termined case, we formulate the constrained optimization $\min_{\{G_{\#}\rho_x=\rho_y\}} E[\rho_x]$. The choice of E also significantly impacts the construction: setting E to be the entropy gives us the piecewise constant reconstruction, while setting E to be the second moment, we recover the classical least-norm solution. We also study the case with regularization: $\min_{\rho_x} D(G_{\#}\rho_x, \rho_y) + \alpha \mathsf{R}[\rho_x]$, and find that the entropy-entropy pair leads to a regularized solution that is defined in a piecewise manner, whereas the W_2-W_2 pair leads to a least-norm solution where W_2 is the 2-Wasserstein metric.

Qin Li
University of Wisconsin-Madison
qinli@math.wisc.edu

Maria Oprea
Cornell University
mao237@cornell.edu

Li Wang
University of Minnesota
wang8818@umn.edu

Yunan Yang
Cornell University
yunan.yang@cornell.edu

MS146

Bayesian Information-Theoretic Approach to Determine Effective Scanning Protocols of Cancer Patients

The aspect of limited temporal data is one of the many challenges when dealing with clinical data. The amount of data that can be practically collected in everyday patients during the therapy is very limited due to the financial cost and the patients burden. This motivates us to transfer the mathematical and computational models to meet the challenges in clinical data, before we use them to guide patient therapy via prediction. In this talk, I will discuss modeling approaches to tackle this problem. For instance, I will discuss a Bayesian information-theoretic approach to determine effective scanning protocols of cancer patients. We propose a modified mutual information function with a temporal penalty term to account for the loss of temporal data. The effectiveness of our framework is demonstrated in determining scanning scheduling for radiotherapy and androgen suppression therapy patients.

Heyrim Cho
Arizona State University
heyrim.cho@asu.edu

Tin Phan
Los Alamos National Laboratory
Theoretical Biology and Biophysics
ttphan@lanl.gov

Tingting Tang
San Diego State University
Dept. of Math & Stat.
ttang2@sdsu.edu

Kathleen Storey, Allison L. Lewis
Lafayette College
storeyk@lafayette.edu, lewisall@lafayette.edu

MS146

Parameter Identifiability of a Hemodynamics Model Using Dimension Reduction and Emulation

Biophysical models have immense potential in providing new tools for the diagnosis and management of disease. This includes cardiovascular monitoring. However, cardiovascular simulations typically require numerous parameters and suffer from limited data. This leads to issues in parameter identifiability, which renders parameters uninterpretable. Assessing identifiability is difficult due to expensive forward solves from the model. Here, we overcome this limitation by using dimension reduction and emulation to expedite computation. We use polynomial chaos expansions (PCEs) on the principal components of a blood flow PDE model to provide a succinct platform for forward model solutions. We use this framework to assess model sensitivity, and follow with a formal identifiability analysis using the profile-likelihood. We show how experimental design choice drastically changes identifiability in the context of blood flow monitoring.

Mitchel J. Colebank
North Carolina State University
mjcolebank@sc.edu

MS146

InVAErt Networks: Model Synthesis and Identifiability Analysis for Scientific Agents

Applications of generative modeling and deep learning in physics-based systems have traditionally focused on building emulators - computationally inexpensive approximations of input-to-output maps. However, the remarkable flexibility of data-driven architectures opens opportunities to broaden their scope to include model inversion and identifiability analysis. We present InVAErt networks, a framework for data-driven analysis and synthesis of parametric physical systems. Through numerical experiments, we demonstrate the framework's versatility across a wide range of problems, including linear systems of equations, spatio-temporal PDEs, and lumped-parameter physiological models. We further introduce an extension for systems with observational noise, enabling the separation of structural from practical identifiability in complex ill-posed inverse problems. Finally, we discuss recent efforts to integrate InVAErt networks with large language model agents for applications in cardiovascular health.

Guoxiang Tong
Stanford University
gxus34@stanford.edu

Carlos Sing Long
Institute for Mathematical and Computational Engineering
Pontificia Universidad Catolica de Chile
casinglo@uc.cl

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

MS147

Efficient Cut-Bayes Computation for Uncertainty

Quantification

Many complex physical systems, including those from the environment, can be modeled as a collection of interconnected modules where each one represents a crucial piece of physics that drives the entire system. For instance, these modules could correspond to atmospheric, ocean, and wave components in an earth system model. One way to enable uncertainty quantification in such multiphysics models is with a cut-Bayesian approach (or more generally, a modular or generalized Bayesian approach). This can help address identifiability and mitigate the possibility of misspecification corrupting inferences. Nonetheless, computation is a challenge for cut Bayesian inference, and in this talk we will discuss a recently published sequential Monte Carlo method that comes with finite sample guarantees. We also discuss connections to inference with AI (e.g., deep learning) models.

Giri Gopalan
Los Alamos National Laboratory
ggopalan@lanl.gov

Joseph Mathews
University of North Carolina at Chapel Hill
Department of Epidemiology
jrm133@unc.edu

MS147

Characterizing Extreme Weather in a Huge Ensemble of Machine Learning Weather Forecasts

Ensemble weather forecasts offer a promising avenue to calculate the statistics and drivers of high-impact extreme weather events. Current ensemble systems presently use up to 100 members, but larger ensembles could enrich the sampling of internal variability. Due to computational constraints, it is infeasible to generate huge ensembles (comprised of 1,000-10,000 members) with traditional, physics-based numerical models. We replace traditional numerical simulations with machine learning to generate hindcasts of huge ensembles. We construct a calibrated ensemble weather system that represents model uncertainty through perturbed-parameter techniques, and it represents initial condition uncertainty through bred vectors, which sample the fastest growing modes of the forecast. Using the Spherical Fourier Neural Operator architecture, we generate a huge ensemble (HENS), with 7,424 members initialized each day of summer 2023, the second hottest summer in at least the last 2000 years. HENS samples events 4s away from the ensemble mean. We use HENS to characterize the spread in atmospheric river intensity and impacts in the Southern Hemisphere. We also apply HENS to ascertain whether tropical cyclones obey conventional or anomalous diffusion subject to large-scale eddies in the general circulation. We use HENS to characterize near-miss counterfactual heatwaves that were plausible in summer 2023.

Ankur Mahesh
Lawrence Berkeley National Laboratory
University of California Berkeley
amahesh@lbl.gov

MS147

A Bayesian Hierarchical Model with a Physics Informed Neural Process for Wind Assimilation

Recently, physics-informed neural networks (PINNs) have gained popularity due to their ability to represent complex physical phenomena parameterized using partial differential equations (PDEs). This is particularly useful in scientific applications where there is a priori knowledge, and incorporating process dynamics provide physically realistic constraints to the solution. However, PINN-based methods rarely account for data or process uncertainty, limiting their use for real-world applications. The notion of including physical information in data-driven methods is well established within the field of statistics, with multiple distinct paradigms that can be used to account for mechanistic information and realistically quantify uncertainty from multiple sources. However, the statistical methods are limited in that realistic dynamics can be computationally difficult to implement. This talk presents a flexible Bayesian hierarchical modeling framework with a neural dynamic process that is able to account for stochastic process discrepancies, observation uncertainty, and parameter uncertainty. The framework is applied to sparsely observed wind data from multiple sources over the Mediterranean Sea, e.g. in situ and satellite measurements, to assimilate the data sources and produce accurate wind field estimates.

Joshua S. North
Lawrence Berkeley National Laboratory
jsnorth@lbl.gov

Myungsoo Yoo
University of Texas at Austin
myungsoo.yoo@austin.utexas.edu

Giri Gopalan
Los Alamos National Laboratory
ggopalan@lanl.gov

Christopher Wikle
Department of Statistics
University of Missouri
wiklec@missouri.edu

MS147

Fast Global-Local Reduced Order Models Using Neural Galerkin-Projections for Improved Tsunami Forecasting

Computationally expensive numerical solutions of partial difference equation models (PDE)s are critical in understanding complex real-world phenomena. However, the utility of these numeric simulations in many applications is limited by the computational cost of running these models over a breadth of initial conditions necessary to characterize the space of feasible solutions. We propose a novel surrogate framework RANDPROM that combines cutting-edge reduced-order approximations of PDEs within a Bayesian hierarchical model that enables accurate and precise predictions of the PDE solutions at initial conditions which are not included in training data. Using simulations of tsunami wave height as an example dataset, we demonstrate the potential for RANDPROM to produce near real-time predictions of wave height which demonstrates the potential of the RANDPROM surrogate framework to contribute to real-world decision making.

John Tipton
Los Alamos National Laboratory, United States
jtipton@lanl.gov

Shane Coffing, Arvind T. Mohan

Los Alamos National Laboratory
sxc@lanl.gov, arvindm@lanl.gov

Darren Engwirda
CSIRO, Australia
d.engwirda@gmail.com

MS148

The Landscape of Sequential Inference for Dynamical Systems

Sequential inference underlies the analysis of stochastic dynamical systems in many areas of applied mathematics and data science. In state-space models, latent states evolve in time and are observed through noisy and partial measurements, which leads to recursive Bayesian inference under modeling and computational constraints. This lecture surveys the principal approaches to sequential inference from a unified mathematical perspective. The presentation begins with the linearGaussian case, where the Kalman filter provides an exact and tractable solution that serves as a reference point. It then examines nonlinear and non-Gaussian models and the approximations they require. Gaussian filters, sequential Monte Carlo methods, and structured variants that exploit conditional independence, RaoBlackwellization, and low-dimensional structure are compared in terms of accuracy, stability, and computational complexity. Deterministic alternatives based on variational approximations and message-passing formulations are also discussed. The lecture concludes with a discussion of fundamental limitations, including dimensional scaling and filter stability, and emphasizes the role of model structure to enable tractable inference. The goal is to clarify the relationships among methods and to identify open challenges in sequential inference for complex dynamical systems.

Petar Djuric
Stony Brook University
petar.djuric@stonybrook.edu

MS148

Fast and Reliable Variational Inference for Uncovering Latent Dynamical Systems

Latent dynamical systems provide a powerful framework for uncovering hidden structure in complex time series, with applications spanning engineering, neuroscience, and beyond. However, inference in these models is often slow and numerically unstable, limiting their scalability in real-world settings. In this talk, I will introduce SING, a variational inference algorithm that leverages natural gradient updates to exploit the geometric structure of latent stochastic differential equation models, enabling fast, stable, and parallelizable inference. I will demonstrate how SING produces accurate estimates of latent trajectories and nonlinear drift functions, including those modeled with interpretable Gaussian process priors. Across synthetic and real datasets, SING consistently improves both convergence speed and inference accuracy, advancing our ability to recover meaningful dynamical structure from high-dimensional, noisy sequential data.

Amber Hu
Standford University
amberhu@stanford.edu

MS149

Physiology-Constrained Uncertainty Quantification for Glucose Dynamics in Type 1 Diabetes

Glucoseinsulin regulation in Type 1 diabetes (T1D) is characterized by nonlinear dynamics, substantial individual variability, and pervasive uncertainty in both physiological processes and sensor measurements. Developing reliable predictive models and decision-support tools requires mathematical frameworks that can rigorously quantify, propagate, and interpret uncertainty while remaining consistent with known biological constraints. This project advances uncertainty quantification (UQ) for nonlinear biological systems by formulating biologically plausible uncertainty representations, establishing constraints that preserve physiological realism, and developing methods for stable uncertainty propagation through dynamical models. The approach integrates traditional differential-equation models with modern computational tools to produce probabilistic forecasts and risk metrics that remain faithful to underlying physiology. The resulting framework provides broadly applicable UQ methodologies for complex biological systems, with T1D serving as a motivating and clinically relevant application.

Yixiang Deng
University of Delaware
yixiangd@udel.edu

MS149

Relative Entropy Methods for Computing Committor Functions

Motivated by challenges arising in molecular simulation, we analyze and develop methods of computing reactive trajectories and committor functions for systems described by the overdamped Langevin dynamics. Our main technical advance is a new loss function that measures the accuracy of approximations to the committor function related to a given chemical reaction or other rare transition event. Our loss admits a simple interpretation in terms of the distribution of reactive trajectories, and it can be computed in practice to compare the accuracies of different approximations of the committor. We also derive a method of calculating committors by direct minimization of the loss via stochastic gradient descent.

Brian Van Koten
Department of Mathematics
University of Massachusetts Amherst
bvankoten@umass.edu

MS149

Mean Square Error Analysis of Stochastic Gradient and Variance-Reduced Sampling Algorithms

This talk presents a rigorous analysis of the mean square error (MSE) for stochastic gradient sampling algorithms applied to underdamped Langevin dynamics under a global convexity assumption. We introduce a novel discrete Poisson equation framework to systematically decompose and bound the time-averaged sampling error. For the primary algorithm, Stochastic Gradient UBU (SG-UBU), we establish an explicit MSE bound and demonstrate that its numerical bias exhibits first-order convergence with respect to the step size h. The leading error coefficient is shown to be directly proportional to the variance of the stochastic gradient. The analysis is extended to variance-reduced algorithms for finite-sum potentials, specifically SVRG-UBU and SAGA-UBU. For these methods, we uncover a distinctive phase transition phenomenon: the convergence rate of

the numerical bias shifts from first-order to second-order as the step size decreases below a critical threshold. Numerical experiments validate these theoretical findings. Finally, our analysis provides a practical empirical criterion for algorithm selection between mini-batch SG-UBU and mini-batch SVRG-UBU to optimize computational cost.

Xuda Ye
Purdue University
ye481@purdue.edu

MS149

Dynamically Adapted Regularization for Autoencoders with Observed SDE dynamics

We introduce two dynamics-adapted regularizers that improve the accuracy and out-of-manifold generalization of autoencoders trained on manifold-valued stochastic dynamics. Assuming access to an SDE solver and to extrinsic estimatessample points on the manifold together with their drift vectors and local covariance matriceswe augment the autoencoder objective with (i) an It-consistency term that enforces agreement between the decoders infinitesimal predictions and the first- and second-order moments prescribed by It's formula, and (ii) a geometry-consistency term that aligns the learned drift with the manifolds tangent space while penalizing normal components and mismatches between decoded covariances and the induced metric. Together these terms couple stochastic moment constraints with manifold geometry, yielding representations whose learned dynamics transfer more reliably across benchmarks. We demonstrate improved generalization on standard examples of manifold SDEs.

Felix Ye
SUNY Albany
xye2@albany.edu

MS150

Reinforcement Learning with Verifiable Rewards: Grpo's Effective Loss, Fixed Point Iteration and Success Amplification

Group Relative Policy Optimization (GRPO) was introduced recently and used successfully to train DeepSeek-R1 models for promoting reasoning capabilities of LLMs using verifiable or binary rewards. We show in this paper that GRPO with verifiable rewards can be written as a KullbackLeibler (KL) regularized contrastive loss, where the contrastive samples are synthetic data sampled from the old policy. The optimal GRPO policy at time n can be expressed explicitly in terms of the binary reward, as well as the first- and second-order statistics of the old policy (at time n-1) and the reference policy . Iterating this scheme, we obtain a sequence of policies for which we can quantify the probability of success as function of n. We show that the probability of success of the policy satisfies a recurrence that converges to a fixed point of a function that depends on the initial probability of success and the regularization parameter beta of the KL regularizer. We show that the success rate corresponding to this fixed point is guaranteed to be larger than the reference policy's success rate, thereby demonstrating that GRPO effectively amplifies the probability of success of the policy.

Youssef Mroueh
IBM Research
mroueh@us.ibm.com

MS150

Simple and Effective Diffusion Language Models

We are at a pivotal moment in language modeling, as autoregressive (AR) LLMs face growing competition from diffusion-based approaches. I will present the Masked Diffusion Language Model (MDLM), which introduces a new training strategy and a simplified Rao-Blackwellized objective that significantly improves discrete diffusion performance, bringing it close to AR models on standard log-likelihood benchmarks. Unlike AR methods, diffusion models offer advantages in long-term planning, controllable generation, and sampling efficiency. Recent releases such as Googles Gemini Diffusion, Inception Labs Mercury Coder, and ByteDances Seed Diffusionwhich is built on the MDLM frameworkunderscore this paradigm shift, signaling that diffusion may soon overtake AR as the dominant framework for language generation.

Subham Sahoo
Computer Science Department
Cornell Tech University
ssahoo@cs.cornell.edu

MS150

Bias and Pitfalls in Aligning Large Language Models: Optimization and Game-Theoretic Perspectives

Large language models (LLMs) are aligned with human preferences through reinforcement learning from human feedback (RLHF). However, alignment is empirically observed to introduce bias, degrade calibration, and only imperfectly reflect diverse human preferences. In this talk, I will present recent progress toward a theoretical foundation of LLM alignment. First, I will show how the current form of RLHF induces a bias we call preference collapse, and how this can be mitigated through a novel regularization term added to the reward function. Second, I will revisit RLHF through the lens of social choice theory and present a game-theoretic perspective, highlighting its violations of classical axiomatic consistency. This talk is based on joint work in arXiv:2405.16455, 2503.10990, 2505.20627, and 2506.12350.

Weijie Su
University of Pennsylvania
suw@wharton.upenn.edu

MS151

Representation Learning for Scientific Datasets Using Disentangled Generative Models

Deep generative models (DGMs) have emerged as powerful tools for learning high-dimensional data distributions and generating realistic samples. However, their latent spaces often entangle multiple independent factors, making it difficult to extract meaningful scientific insight. I will present a disentangled generative modeling framework that introduces auxiliary physical variables to guide the latent space, aligning learned representations with known factors while retaining generative flexibility. This approach not only enhances interpretability but also improves data efficiency by leveraging partial prior knowledge to constrain learning. I will illustrate the method in two astrophysical settings: galaxy images generated with Gal-

Sim, where the model recovers key morphological features, and cosmic microwave background (CMB) halo maps from the thermal SunyaevZeldovich effect, where it disentangles halo mass and concentration while revealing outlier populations. More broadly, this auxiliary-variable-guided framework provides a general pathway for data-efficient representation learning in complex scientific domains.

Arkaprabha Ganguli
Argonne National Laboratory
Mathematics and Computer Science Division
aganguli@anl.gov

MS151

Efficient Estimation of Information Theoretic Quantities and Their Application in Generative Models

Information theoretic quantities play a central role in machine learning. The recent surge in the complexity of data and models has increased the demand for accurate estimation of these quantities. However, as the dimension grows the estimation presents significant challenges, with existing methods struggling already in relatively low dimensions. In this talk, we introduce REMEDI for efficient and accurate estimation of differential entropy, a fundamental information theoretic quantity. The approach combines the minimization of the cross-entropy for simple, adaptive base models and the estimation of their deviation, in terms of the relative entropy, from the data density. Our approach demonstrates improvement across a broad spectrum of entropy estimation tasks. We illustrate how the framework can be naturally extended to information theoretic supervised learning models, with a specific focus on the Information Bottleneck approach. In addition, we explore a natural connection between REMEDI and generative modeling using rejection sampling and Langevin dynamics.

Anirban Samaddar
Argonne National Laboratory
asamaddar@anl.gov

MS151

Belted and Ensembled Neural Network for Linear and Nonlinear Sufficient Dimension Reduction

We introduce a unified, flexible, and easy-to-implement framework of sufficient dimension reduction that can accommodate both linear and nonlinear dimension reduction, and both the conditional distribution and the conditional mean as the targets of estimation. This unified framework is achieved by a specially structured neural network – the Belted and Ensembled Neural Network (BENN) – that consists of a narrow latent layer, which we call the belt, and a family of transformations of the response, which we call the ensemble. By strategically placing the belt at different layers of the neural network, we can achieve linear or nonlinear sufficient dimension reduction, and by choosing the appropriate transformation families, we can achieve dimension reduction for the conditional distribution or the conditional mean. Moreover, thanks to the advantage of the neural network, the method is very fast to compute, overcoming a computation bottleneck of the traditional sufficient dimension reduction estimators, which involves the inversion of a matrix of dimension either p or n. We develop the algorithm and convergence rate of our method, compare it with existing sufficient dimension reduction methods, and apply it to two data examples.

Yin Tang
University of Kentucky
yin.tang@uky.edu

Bing Li
Pennsylvania State University
bxl9@psu.edu

MS152

On Some Foundational Issues in Uncertainty Control

The remarkable success of closed-loop control in mitigating the effect of uncertainty on a systems performance has undoubtedly enabled much of the technological world around us. While the mitigation of uncertainty has been at the heart of control theory since its inception, explicit control of uncertainty is a relatively recent development that has garnered much attention. In this, a main object of study is the Liouville (continuity) equation – the PDE governing the evolution of the probability distribution of the state of a dynamical system. While it was widely believed that the basic question of controllability of the Liouville equation had been resolved, it escaped the communitys attention for almost two decades that early investigations on the subject came short of providing a satisfactory answer, even for linear systems. In this talk, we revisit and address this topic and develop a theory for Collective Steering, the endeavor to shepherd an ensemble of dynamical systems between desired configurations using a common feedback law. Our investigation sheds light on a topological obstruction at the heart of the issue that limits the ability to design feedback control laws that are globally continuous with respect to the specifications. Along the way, we touch upon an elegant geometric framework at the intersection of optimal transport, geometric hydrodynamics, and quantum mechanics.

Mahmoud Abdelgalil
UC San Diego
mabdelgalil@ucsd.edu

MS152

Wasserstein Least Squares Regression

We present Wasserstein Least Squares Regression (WLSR), a model that canonically extends least squares regression in the presence of vector-valued covariates and distribution-valued responses. Unlike competing proposals, which focus on the linear structure in the space of probability measures, ours works directly with the functional form of linear regression. In this talk, we will delve into the geometry of WLSR and draw connections with the theory of Mixed linear models.

Uriel Martinez Leon
NYU
um2065@nyu.edu

MS152

Error Analysis of Triangular Optimal Transport Maps for Nonlinear Filtering

The stochastic filtering problem seeks to recover the state of a dynamical system from sequential, noisy observations. Classical approaches such as the Kalman and particle fil-

ters often suffer from instability, sampling inefficiency, or intractability in non-linear, non-Gaussian regimes. The Optimal Transport Filter (OTF) circumvents these issues by reformulating the Bayesian update as a conditional optimal transport (COT) problem: the posterior distribution at each step is obtained by pushing forward prior samples through a learned approximation of the COT map. In this paper, we: 1) Extend recent optimal transport map estimation rates to the conditional setting of COT. 2) Introduce a conditional extension tailored to Bayesian inversion in sequential settings to provide the first statistically complete error analysis of the OTF. 3) Investigate and derive sufficient regularity assumptions on the dynamics and observation models defining the filtering problem to obtain quantitative guarantees on the fidelity of posterior approximation, thereby placing the OTF within a principled theoretical framework. 4) Introduce regularization techniques to close the gap between the theoretical and numerical algorithms. These results sharpen the understanding of transport-based inference, bridge connections between stochastic filtering and conditional optimal transport, and inform the design of provably reliable algorithms for high-dimensional Bayesian estimation.

Mohammad Al-Jarrah, Bamdad Hosseini, Niyizhen Jin, Michele Martino
University of Washington
mohd9485@uw.edu, bamdadh@uw.edu, njin2@uw.edu, michem29@uw.edu

Amirhossein Taghvaei
University of Washington Seattle
amirtag@uw.edu

MS153

Advances in Surrogate-Accelerated Calibration for Quasi-Periodic Data

Simulating climate phenomena like the quasi-biennial oscillation (QBO) is challenging for Earth system models due to uncertain subgrid-scale processes like convective gravity waves. We present a novel workflow that improves the QBOs representation in the Energy Exascale Earth System Model (E3SM) by using machine learning (ML) to systematically calibrate unresolved physics, replacing subjective tuning. Our methodology begins with the fundamental frequency model (FFM), a domain-informed statistical method we developed to compress high-dimensional, spatio-temporal wind field data. The FFM isolates the QBO signal into two physically interpretable quantities of interest: period and amplitude. We then develop an efficient Gaussian process surrogate, a probabilistic approximation of the full physics model that rigorously links E3SM to ML. This surrogate enables rapid, low-cost exploration of the parameter space. Our framework also facilitates model-predictive decision-making. By employing multi-objective optimization, we navigate the fundamental trade-off between achieving the correct QBO period and amplitude, quantifying a Pareto frontier of optimal solutions instead of a single parameter set. This integrated workflow has improved QBO simulation fidelity in E3SMv3 and provides a transferable template for incorporating ML to model quasi-periodic processes in complex Earth system models.

Luis Damiano, Khachik Sargsyan
Sandia National Laboratories
ladamia@sandia.gov, ksargsy@sandia.gov

Bert J. Debusschere
Sandia National Laboratories, U.S.
bjdebus@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS154

Automatic RKHS for Learning Convolution Kernels

Learning convolution kernels in operators from data arises in numerous applications and represents an ill-posed inverse problem of broad interest. With scant prior information, kernel methods offer a natural nonparametric approach with regularization. However, a major challenge is to select a proper reproducing kernel, especially as operators and data vary. We show that the input data and convolution operator themselves induce an automatic, data-adaptive RKHS (DA-RKHS), obviating manual kernel selection. In particular, when the observation data is discrete and finite, there is a finite set of automatic basis functions sufficient to represent the estimators in the DA-RKHS, including the minimal-norm least-squares, Tikhonov, and conjugate-gradient estimators. We develop both Tikhonov and scalable iterative, hybrid, and divide-and-conquer algorithms using the automatic basis functions. Numerical experiments on integral, nonlocal, and aggregation operators confirm that our automatic RKHS regularization consistently outperforms standard ridge regression and Gaussian process methods with preselected kernels.

Fei Lu
Johns Hopkins University
flu15@jh.edu

Haibo Li
The University of Melbourne, Australia
haibo.li@unimelb.edu.au

MS154

The Measure Theoretic Koopman Operator

The Stochastic Koopman (SKO) operator has been associated in literature with dynamical systems that exhibit stochastic behaviour. However, the SKO can only recover the expectation of an observable along the trajectory and requires tracking data for its computation. In this work, the Distributional Koopman Operator (DKO) is introduced as a novel framework for analyzing stochastic dynamical systems where only aggregate distribution data is available, avoiding the need for particle tracking. Our DKO framework extends the Koopman operator to act on observables of probability distributions, leveraging the transfer operator to propagate these distributions forward in time. We analyse the properties of this new operator, and show how one can recover information about higher order moments from this formulation. Lastly, we design an algorithm to compute the DKO from data.

Maria Oprea
Cornell University

mao237@cornell.edu

MS154

Randomized Methods for Operator Learning

Operator learning is a technique for approximating operators between two infinite dimensional Banach spaces of functions using machine learning. One approach is a kernel-based method, which involves the use of a kernel function between pairs of data points. However, for large datasets, this becomes quite computationally expensive. We propose a regularized random Fourier feature method to learn operators in the regime of noisy data. The random feature method is an approximation to a kernel method that uses randomly sampled feature maps and produces comparable results with similar test errors but with significantly lower training times. We provide a detailed study of our proposed method with comprehensive numerical results.

Tracy Yu
UCLA
tracy@math.ucla.edu

MS154

Approximation of the Electrical Impedance Tomography Inversion Operator by Neural Operators

We present the noise-robust neural operator approximation of the solution map for Caldern's inverse conductivity problem. In this continuum model of electrical impedance tomography (EIT), the boundary measurements are realized as a noisy perturbation of the Neumann-to-Dirichlet map's integral kernel. The theoretical analysis proceeds by extending the domain of the inversion operator to a Hilbert space of kernel functions making use of the Benyamin-Lindenstrauss theorem. The resulting extension shares the same stability properties as the original inverse map, but is now amenable to neural operator approximation. The methodology for EIT exemplifies a broader framework for addressing nonlinear inverse problems with a noise-aware operator learning framework. We give a universal approximation theorem with certain neural operators with explicit expression rates accounting for the modulus of continuity of the (inversion) operator to be approximated. Joint research with N. Kovachki, M. Lassas and N. Nelsen.

Maarten de Hoop
Rice University
mvd2@rice.edu

MS155

Ensemble Score Filter for Data Assimilation of Two-Phase Flow Models in Porous Media

Numerical modeling and simulation of two-phase flow in porous media is challenging due to the uncertainties in key parameters, such as permeability. To address these challenges, we propose a computational framework by utilizing the novel Ensemble Score Filter (EnSF) to enhance the accuracy of state estimation for two-phase flow systems in porous media. The forward simulation of the two-phase flow model is implemented using a mixed finite element method, which ensures accurate approximation of the pressure, the velocity, and the saturation. The EnSF leverages score-based diffusion models to approximate filtering distributions efficiently, avoiding the computational expense of neural network-based methods. By incorporating a closed-form score approximation and an analytical update mechanism, the EnSF overcomes degeneracy issues and handles high-dimensional nonlinear filtering with minimal computational overhead. Numerical experiments demonstrate the capabilities of EnSF in scenarios with uncertain permeability and incomplete observational data.

Ruoyu Hu
Florida State University
rhu2@fsu.edu

MS155

Ensemble Kalman-Bucy Smoother and Applications

In this talk, we will present the mathematical formula for the ensemble Kalman-Bucy smoother, a continuous-time data assimilation method. We will show the stochastic version of the smoother. Then we will apply this method to various problems, including causal inference via uncertainty quantification and data assimilation, model identification and parameter estimation from partial observations, and state estimation for extreme events.

Zhang Jiang
University of Wisconsin-Madison
zjiang368@wisc.edu

Nan Chen, Nan Chen
University of Wisconsin - Madison
chennan@math.wisc.edu, chennan@math.wisc.edu

Marios Andreou
University of Wisconsin-Madison
mandreou@wisc.edu

MS155

Operator-Free Data Assimilation for Complex Measurement Processes

Data assimilation traditionally relies on a measurement operator, or forward model, to link model states to observations. In many applications, the measurement process is occluded by bias or is imperfectly known. We present a suite of operator-free strategies for estimating latent states directly from observations when the forward model is unknown. First, we describe deterministic surrogate observation models in which neural networks are trained to approximate the forward mapping from state space to observation space, providing flexiblebut not necessarily operator-freeestimates. Second, we introduce a strategy that leverages kernel mean embeddings to represent conditional distributions in a reproducing kernel Hilbert space, enabling nonparametric inference without explicit forward operators. Third, we explore probabilistic generative methods based on variational autoencoders, which learn latent representations that capture the joint structure of states and observations and can be constructed to accommodate measurement uncertainty. To demonstrate the scalability of these approaches, we show how they can be integrated with current filtering algorithms for high-dimensional geophysical systems, including numerical weather prediction, sea ice state estimation, and paleoclimate reconstruction. Results highlight trade-offs in accuracy, cost, and UQ across the three strategies, illustrating a path toward data assimilation for next-generation Earth System models.

Jonathan Poterjoy
UCAR

poterjoy@umd.edu

Raelynn Wonnacott, Henry Santer, Michael Evans
University of Maryland
raelynnw@umd.edu, hsanter@umd.edu, mnevans@umd.edu

MS155

Neural Ensemble Kalman Filter: Data Assimilation for Compressible Flows with Shocks

Data assimilation (DA) in compressible flows with shocks is challenging. Standard DA methods, such as the ensemble Kalman filter (EnKF), often generate oscillatory and nonphysical solutions near shocks. We introduce a latent-space formulation of the EnKF, termed the "neural EnKF," which enables robust DA in the presence of shocks. The neural EnKF operates on an ensemble of deep neural networks, each parameterizing a realization of the shocked flow field. The networks share the same architecture and are trained successively using a nearest-neighbor chain, which ensures that shocks in physical space correspond to smooth variations in neural parameter space. Observations in physical space (e.g., pressure) are then assimilated to update the network weights, yielding an analysis ensemble that represents physically consistent shocked flows informed by the data. We present systematic numerical DA experiments with the inviscid Burgers' equation and Sod's shock tube to demonstrate the robustness, efficiency, and accuracy of the neural EnKF, highlighting its ability to preserve sharp shock structures while avoiding spurious oscillations.

Xuhui Zhou, Lorenzo Beronilla
University of California San Diego
xuz067@ucsd.edu, lberonilla@ucsd.edu

Michael Sleeman
California Institute of Technology
msleeman@caltech.edu

Matthias Morzfeld
Scripps Institution of Oceanograpghy
University of California San Diego
matti@ucsd.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

Tamer Zaki
Johns Hopkins University
t.zaki@jhu.edu

MS156

Approximation of Risk-Averse Optimal Feedback Control

The challenge of constructing feedback control laws for risk-averse optimal control of partial differential equations with random coefficients is addressed. The control objective composes a tracking-type cost with the nonlinear entropic risk measure. A sequential quadratic programming scheme is derived that iteratively solves linear quadratic subproblems obtained through second-order Taylor expansions of the objective functional, with each subproblem re-centered at the previous iterate. It is shown that this method converges locally quadratically to the unique risk-averse optimal control.

Philipp A. Guth
RICAM, Austrian Academy of Sciences
philipp.guth@ricam.oeaw.ac.at

Karl Kunisch
University of Graz and Radon Institute, Austrian Academy of
Science, Austria
karl.kunisch@uni-graz.at

MS156

New Optimality Conditions and an Indirect Shooting Method for Stochastic Optimal Control

Efficiently solving nonlinear stochastic optimal control problems remains a challenge with many applications. Existing optimality conditions typically rely on forward-backward stochastic differential equations (FBSDEs), which are difficult to use in algorithms. In this talk, I will present new optimality conditions for stochastic optimal control, derived using rough path theory. This Pontryagin Maximum Principle uses the same Hamiltonian as in the deterministic setting, while avoiding FBSDEs. This result unlocks the first indirect shooting method for stochastic optimal control, which only searches over the initial adjoint state and converges an order of magnitude faster than a direct method in numerical experiments. I will conclude with applications to driving at the limits of handling and future directions towards robust and efficient uncertainty-aware control.

Thomas Lew
Toyota Research Institute
thomas.lew@tri.global

MS156

Central Limit Theorems for Sample Average Approximations in Stochastic Optimal Control

We establish central limit theorems for the Sample Average Approximation (SAA) method in discrete-time, finite-horizon Stochastic Optimal Control. Using the dynamic programming principle and backward induction, we characterize the limiting distributions of the SAA value functions. The asymptotic variance at each stage decomposes into two components: a current-stage variance arising from immediate randomness, and a propagated future variance accumulated from subsequent stages. This decomposition clarifies how statistical uncertainty propagates backward through time. Our derivation relies on a stochastic equicontinuity condition, for which we provide sufficient conditions. We illustrate the variance decomposition using the classical Linear Quadratic Regulator (LQR) problem. Although its unbounded state and control spaces violate the compactness assumptions of our framework, the LQR setting enables explicit computation and visualization of both variance components.

Johannes Milz
Georgia Institute of Technology
H. Milton Stewart School of Industrial & Systems Engineering
johannes.milz@isye.gatech.edu

Alexander Shapiro
Georgia Institute of Technology

School of Industrial and Systems Engineering
ashapiro@isye.gatech.edu

MS156

Neural Network Approaches for High-Dimensional Optimal Control and Transport Problems

We present neural network approaches for high dimensional optimal control and optimal transport problems. The Neural-HJB method approximates optimal policies, enabling real-time control while alleviating the computational burden of traditional approaches. Similarly, the COT-Flow method applies DNNs to solve dynamic conditional optimal transport problems, which is critical for applications like sampling and density estimation in Bayesian inference. Empirical results highlight the effectiveness of both methods, showcasing their performance compared to state-of-the-art alternatives.

Deepanshu Verma, Deepanshu Verma
Clemson University
dverma@clemson.edu, dverma@clemson.edu

MS157

Transport Maps for Scalable Uncertainty Quantification of Fine Resolution Remote-Sensing Data Products

Remote sensing data are providing valuable information for science and applications at fine spatial and temporal resolution. In environmental applications, these utility of these science data products is often aided with complementary information from multiple instruments, as well as from geophysical models. This added utility can come at added computational cost with additional complexity for uncertainty quantification (UQ) due to variable fidelity of input uncertainty and spatio-temporal correlation. We present a framework for UQ for these derived products that combines generative flexibility of transport maps (TMs) with scalable spatial statistical methods for dense geophysical fields. The presentation will highlight the implementation for estimates of atmospheric aerosols arising from a remote sensing inverse problem for NASAs Multi-Angle Imager for Aerosols mission. The mission will deliver multivariate aerosol products for 1-km footprints over a collection of target areas across global land areas.

Jonathan Hobbs, Otto Lamminpää, Maggie Johnson
Jet Propulsion Laboratory
jonathan.m.hobbs@jpl.nasa.gov,
otto.m.lamminpaa@jpl.nasa.gov,
maggie.johnson@jpl.nasa.gov

Ricardo Baptista
University of Toronto
r.baptista@utoronto.ca

Olga Kalashnikova
Jet Propulsion Laboratory
olga.kalashnikova@jpl.nasa.gov

Meredith Franklin
University of Toronto
meredith.franklin@utoronto.ca

Gregory Halverson
Jet Propulsion Laboratory
gregory.h.halverson@jpl.nasa.gov

Kerry Cawse-Nicholson
Jet Propulsion Laboratory
California Institute of Technology
kerry-anne.cawse-nicholson@jpl.nasa.gov

MS157

A Bayesian Point Process Tree Model for High-Dimensional Data

Spatial point patterns arise in many scientific applications from sociology to climate science and to biology. Statistical analysis of spatial point patterns often needs to deal with heterogeneous structures in intensity function estimation. Likelihood based inference for spatial point process models often needs to deal with computational and statistical challenges due to the evaluation of multidimensional integral of intensity function and nonstationarity modeling of intensity functions over possibly irregular domain. In this talk, I will introduce a Bayesian Poisson process tree (PPT) model that allows exact Bayesian inference without relying on integral approximation. In particular, the intensity function model is modeled via a tree structure that is endowed with adaptive shrinkage prior to local features. To allow efficient Bayesian inference, we develop a sequential Monte Carlo algorithm to allow efficient tree learning. The proposed method is demonstrated with various simulation studies and a real application in high-dimensional covariance space.

Pulong Ma
Department of Statistics
Iowa State University
plma@iastate.edu

MS158

Uqpy Version 4.2: A Modular Python Framework for Uncertainty Quantification and Scientific Machine Learning

UQpy is an open-source Python library for conducting a wide range of uncertainty-quantification (UQ) studies on physical and mathematical systems. In its previous release, it comprises nine modules covering complementary aspects of UQfrom probability distributions and sampling to surrogate modeling of complex systems and state-of-the-art dimension-reduction methods such as Grassmann manifold projections. It adopts an enhanced object-oriented architecture for a simpler user/developer experience, while continuous integration enforces code quality and automated tests validate scientific results. System-agnostic packaging via PyPI, conda-forge, and Docker removes platform constraints and eases dissemination. UQpy Version 4.2 introduces a new Scientific Machine Learning Module (SciML) module that extends UQpy into machine-learning workflows. Built on PyTorch, it provides custom layers, neural networks, and trainers compatible with?torch 2.2.2, designed for plug-and-play use within existing PyTorch code. The module effectively extends PyTorch with UQ capabilities for both neural networks and?neural operators, including Bayesian implementations of common layers and the associated training procedures to learn their parameters. A flexible, object-oriented interface lets users define?Fourier Neural Operators (FNOs),?Deep Operator Networks (DeepONets), and?U-shaped neural operators, enabling robust UQ in modern scientific ML models.

Dimitris Giovanis, Ponkrshnan Thiagarajan, Michael D. Shields
Johns Hopkins University

dgiovan1@jhu.edu, ponkrshnan@jhu.edu, michael.shields@jhu.edu

MS158

Dakota: Software for Black-box UQ and Optimization

Dakota is an open-source and widely used software package for performing uncertainty quantification, optimization, sensitivity analysis, and model calibration on black box computational simulations. Begun as an internal research project at Sandia National Laboratories almost 30 years ago, we estimate it now has thousands of users worldwide in academia, government, and industry. This talk will provide a brief introduction to and review of Dakotas history and capabilities. Particular attention is given to recent developments in algorithms and in Dakotas general usability, especially in multifidelity and multilevel sampling methods. In addition, nascent efforts to refresh Dakotas architecture and user interface will be described. SAND2021-15837A Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DENA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Adam Stephens
Sandia National Laboratories
jasteph@sandia.gov

MS159

Approximation of Differential Entropy in Bayesian Optimal Experimental Design

Bayesian optimal experimental design provides a framework for selecting experiments that maximize obtained information. We focus on estimating the expected information gain in the setting where the differential entropy of the likelihood is either independent of the design or can be evaluated explicitly. This reduces the problem to maximum entropy estimation. Our study is motivated by large-scale inference problems, such as inverse problems, where the computational cost is dominated by expensive likelihood evaluations. We propose a computational approach in which the evidence density is approximated by a Monte Carlo or quasi-Monte Carlo surrogate, while the differential entropy is evaluated using standard methods without additional likelihood evaluations. We prove that this strategy achieves convergence rates that are comparable to state-of-the-art methods for full expected information gain estimation, particularly when the cost of entropy evaluation is negligible. Our approach relies only on mild smoothness of the forward map and avoids strong technical assumptions. We present numerical experiments that support our theoretical findings.

Chuntao Chen, Tapio Helin
LUT University
gd.sophie@gmail.com, Tapio.Helin@lut.fi

Nuutti Hyvönen
Aalto Univeristy
nuutti.hyvonen@aalto.fi

Yuya Suzuki
Aalto University
yuya.suzuki@aalto.fi

MS159

Optimal Experimental Design of Bayesian Inverse Problems with Multi Sensor Data: Formulation and Algorithms

The sensor placement problem is among the most fundamental questions in optimal experimental design for Bayesian inverse problems. In the common setup, each sensor is identical in terms of quality and cost, and thus the problem is typically constrained by just the total number of sensors. We explore the possibility of different sensor types with varying costs. The resulting optimal experiment design problem is a knapsack constrained optimization problem. We formulate the D-optimal design problem and present a randomized algorithm for solving the it in the case of a linear inverse problem. In practice, this algorithm has noticeable speed ups against the deterministic greedy algorithm at the expense of a reasonable decrease in the solution quality. For nonlinear inverse problems, we utilize a Bayesian Approximation Error (BAE) approach.

Steven Maio
NCSU
smaio@ncsu.edu

MS159

Batch-based Optimal Experimental Design in Bayesian Inverse Problems

Optimal experimental design (OED) aims to optimize the measurement setting of a problem so that the information obtained from observed variables is maximized. In this presentation, we adopt the Bayesian approach to OED and consider batch-based OED for inverse problems. The aim in batch-based OED is to optimize several design variables simultaneously. Optimizing multiple design variables at once is beneficial in situations where multiple measurements can be carried out at once or the next measurement is not immediately available. In this presentation, we present a numerical method for the batch-based OED optimization problem and also consider related theoretical aspects.

Sofia Mäkinen, Tapio Helin
LUT University
sofia.makinen@lut.fi, Tapio.Helin@lut.fi

Andrew Duncan
Imperial College London
a.duncan@imperial.ac.uk

MS159

Non-Intrusive Optimal Experimental Design for Large-Scale Nonlinear Bayesian Inverse Problems Using a Bayesian Approximation Error Approach

We consider optimal experimental design (OED) for nonlinear inverse problems within the Bayesian framework. We propose an approach based on computing an off-line linearisation to the associated forward problem and then optimizing the experimental design. Replacing an accurate model with a linear surrogate, while justified for certain problems, typically leads to incorrect posteriors and sub-optimal designs if the model discrepancy is ignored. We propose the

use of Bayesian approximation error (BAE) approach to account for this while formulating an A-optimal design objective for sensor selection that is aware of the model error. We exploit a key invariance property which leads to a particularly straightforward approach which naturally accommodates marginal OED and also non-differentiable forward models. We demonstrate the effectiveness of our method on several PDE-governed examples. This is joint work with K. Koval

Ruanui Nicholson
Department of Engineering Science
University of Auckland
ruanui.nicholson@auckland.ac.nz

MS160

Projected Stein Variational Inference for Large-Scale Bayesian Inverse Problems Using Derivative Informed Neural Operators

Bayesian inversion of parametric PDEs is a powerful tool for quantifying uncertainty in scientific models, but its application is often limited by the high cost of repeated PDE solves in high-dimensional settings. We present a surrogate-assisted projected Stein variational gradient descent (pSVGD) framework that addresses these challenges. Our approach integrates data-driven model reduction with a derivative-informed neural operator that learns the PDE solution operator in reduced space, amortizing the cost of PDE solves across inference iterations. By combining parameter and solution projections with accurate, derivative-informed surrogate training, we obtain efficient and reliable score function evaluations for particle updates. Theoretical error bounds provide control of surrogate accuracy, while numerical results demonstrate substantial computational savings without loss of inference fidelity. This framework opens the door to fast and scalable Bayesian inference for complex PDE models with high-dimensional parameters.

Benjamin A. Burns, Yuan Qiu, Peng Chen
Georgia Institute of Technology
bburns46@gatech.edu, yuan.qiu@gatech.edu, pchen402@gatech.edu

MS160

Goal-Oriented Model Construction for Multifidelity Approximate Control Variate Estimation

Multifidelity methods for uncertainty quantification are critical for achieving optimal bounds on estimator variance under resource constraints. When a variety of low-fidelity models are available, approaches such as multifidelity Monte Carlo (MFMC) and approximate control variates (ACV) often achieve orders of magnitude lower variance than standard Monte Carlo estimation for the same computational budget. However, most low-fidelity model ensembles are subject to additional resource constraints largely ignored by this process. This talk considers the use of projection-based reduced order models (ROMs) in a multifidelity ACV workflow, where the information cost of training these ROMs is included in the resulting optimization. By balancing training and sampling costs of the model ensemble, the proposed method automates the construction of ROMs according to their predicted utility in multifidelity estimation. Results on exemplars from high-speed flows, including the HiFIRE-1 test vehicle, show reduced variance of greater than one order-of-magnitude from this procedure, despite using imperfect estimates of unconstructed ROM statistics.

Anthony Gruber
Sandia National Laboratories
adgrube@sandia.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Patrick J. Blonigan
Sandia National Laboratories
pblonig@sandia.gov

Eric Parish
Sandia National Lab
eric.parish9@gmail.com; ejparis@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

MS160

Bayesian Optimization for Efficient Assimilation and Uncertainty Quantification of Wall-Pressure Measurements in Hypersonic Flow over a Cone

Data assimilation involves solving an inverse problem to identify the state of a dynamical system from experimental measurements. Solving such inverse problems for complex phenomena, such as transitional hypersonic flows, is computationally costly using direct numerical simulations (DNS). Surrogate models could alleviate the computational costs. However, data generation to train the surrogate models still involves performing expensive DNS. In this work, we implement an active learning (AL) framework and a Bayesian optimization (BO) framework to reduce the computational cost of the data assimilation problem to identify the inflow disturbances for a Mach 6, transitional boundary layer flow over a cone from experimental measurements. The goal of the AL framework is to learn the surrogate model at a given accuracy with the least number of training data points. The AL framework selects the most informative points to train the surrogate model at every iteration. As such, the number of DNS required to generate the training data is reduced, resulting in a significant reduction of the computational cost. Once the surrogate model is trained using the AL framework, it can be used to solve the inverse problems of interest. In contrast, the goal of the BO framework is to solve a given inverse problem in the smallest number of iterations possible. We demonstrate the computational benefits of the proposed methods and quantify uncertainties in the inversely identified inflow disturbances.

Ponkrshnan Thiagarajan, Michael D. Shields, Tamer A. Zaki
Johns Hopkins University
ponkrshnan@jhu.edu, michael.shields@jhu.edu, t.zaki@jhu.edu

MS160

UQ for Synthetic Time Series Generated from Multimodal Auxiliary Data Streams

Monitoring the health and functioning of physical assets is essential to many industrial and scientific endeavors, but

often requires examining data that can be difficult to obtain or potentially cost-prohibitive. Recently, there has been a growth in methods that utilize inexpensive, continuous monitoring tools to obtain alternative data sources as a proxy for such expensive data. By utilizing modern ai/ml approaches, experts seek to create models that can infer and generate synthetic data on key quantities of interest from the fusion of inexpensive, auxiliary sources. However, these black-box models increase the risk of inaccurate quality assessment and inferences without appropriate uncertainty guarantees. In this talk, we discuss a UQ approach for synthetic time series data generated from auxiliary measurements of a physical ratcheting device. We utilize a conformal predictions approach with our black-box generative model to provide robust UQ for relevant quantities of interest. Comparisons with experimental results are discussed.

Tian Yu Yen
Sandia National Laboratories
tyen@sandia.gov

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Amanda Massey, Jason Lund, Tiernan Casey
Sandia National Laboratories
almasse@sandia.gov, jmlund@sandia.gov, tcasey@sandia.gov

MS161

Learned Hemodynamic Coupling Inference in Resting-State Functional MRI

Functional magnetic resonance imaging (fMRI) provides an indirect measurement of neuronal activity via hemodynamic responses that vary across brain regions and individuals. Ignoring this hemodynamic variability can bias downstream connectivity estimates. Furthermore, the hemodynamic parameters themselves may serve as important imaging biomarkers. Estimating spatially varying hemodynamics from resting-state fMRI (rsfMRI) is therefore an important but challenging blind inverse problem, since both the latent neural activity and the hemodynamic coupling are unknown. In this work, we propose a methodology for inferring hemodynamic coupling on the cortical surface from rsfMRI. Our approach avoids the highly unstable joint recovery of neural activity and hemodynamics by marginalizing out the latent neural signal and basing inference on the resulting marginal likelihood. To enable scalable, high-resolution estimation, we employ a deep neural network combined with conditional normalizing flows to accurately approximate this intractable marginal likelihood, while enforcing spatial coherence through priors defined on the cortical surface that admit sparse representations. Uncertainty in the hemodynamic estimates is quantified via a double-bootstrap procedure. The proposed approach is extensively validated using synthetic data and real fMRI datasets, demonstrating clear improvements over current methods for hemodynamic estimation and downstream connectivity analysis.

William Consagra
University of South Carolina
consagra@mailbox.sc.edu

MS161

Navigating the Optogenetic Landscape: How Theory and Experiment Reveal Rules of Recurrent Cortical Computation

Most neocortical synapses are recurrent, yet the functional logic governing when recurrence promotes cooperation versus competition among neuronal ensembles is poorly understood. Multiphoton holographic optogenetics offers a novel approach to this question, enabling precise activation of defined neuron ensembles in primary visual cortex (V1) without confounding external stimuli. However, these experiments lead to a vast parameter space, as ensemble size, spatial configuration, and tuning composition can all be varied independently. Here, we combined targeted optogenetic perturbations with computational modeling in an iterative framework. Experiments show that spatial arrangement and stimulus preference of activated ensembles shape network responses. Compact, co-tuned ensembles recruit nearby iso-tuned neurons while suppressing others. These observations motivated a computational model incorporating local excitatory connectivity and selective inhibition. The model recapitulated findings and generated new predictions about the trade-off between excitatory and suppressive pathways, which were tested experimentally. Combining in vivo perturbation and in silico modeling clarifies how recurrent cortical circuits selectively amplify or suppress activity depending on the precise pattern of co-active neurons, demonstrating how theory-experiment dialogue can extract organizing principles from high-dimensional experimental landscapes.

Gregory Handy
University of Minnesota
ghandy@umn.edu

MS162

Joint Parameter and State Dimension Reduction for Bayesian Ice Sheet Inverse Problems

Solving large-scale Bayesian inverse problems governed by complex PDEs suffers from the twin difficulties of the high-dimensionality of the uncertain input parameters and the cost of solving computationally expensive forward models. We present a framework for joint parameter and state dimension reduction that mitigates these computational challenges. To reduce the parameter dimension we leverage local sensitivity of the data to parameters, the smoothing properties of the forward model, the fact that the data contain limited information about the infinite-dimensional parameter field, and the covariance structure of the prior. The method identifies a likelihood-informed parameter subspace that captures where the change from prior to posterior is most significant. For the state dimension reduction, we employ a proper orthogonal decomposition to reduce the linear part of the forward solve combined with the empirical interpolation method to approximate the nonlinear term in the forward model. We apply this joint input parameter and state space reduction strategy to a Bayesian ice sheet inverse problem in which we recover the basal sliding coefficient using partial surface velocity data. The input space is reduced from 135 to 15 dimensions, and the state space is reduced to 15 dimensions from its original size of 6179. Despite this drastic dimension reduction we see excellent agreement between the MAP point for the recovered basal friction coefficient and the true parameter.

Susan Minkoff
Brookhaven National Labs
Brookhaven National Labs
sminkoff@bnl.gov

Preskella Mrad
University of Texas at Dallas
preskella.mrad@utdallas.edu

Noémi Petra
University of California Merced
npetra@ucmerced.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

Tiangang Cui
Monash University
tiangang.cui@sydney.edu.au

Ki-Tae Kim
University of California Merced
kkim107@ucmerced.edu

MS162

Exploring Inputs and Suitability of a Model for Prediction from Deep Learning Models of Post Fire Debris Flows

In this talk we will discuss recent work on developing deep learning-based surrogates for post fire debris flow models and their use in new sites for which training data is sparse. Topography, post-burn conditions, and rainfall are the primary drivers of such flows. We focus here on strategies for effective prediction at new sites using ideas from optimal transport and spectral clustering.

Abani Patra
Tufts University
abani.patra@tufts.edu

MS162

VPPE: Application of Scaled Vecchia Approximations to Parallel Partial Emulation

Computer models or simulators are widely used across scientific fields, but are computationally expensive limiting their use to explore possible scenarios/outcomes. Gaussian process emulators are statistical surrogates that can rapidly approximate the outputs of computer models at untested inputs and enable uncertainty quantification studies. The parallel partial emulation (PPE) was developed to model simulators with vector-valued outputs. While the PPE is adept at fitting simulator data with multidimensional outputs, the time to fit the PPE increases quickly as the number of training runs increases. The Scaled Vecchia approximation, a fast approximation to multivariate Gaussian likelihoods, makes fitting Gaussian process emulators with large training datasets tractable. Here we introduce the Vecchia Parallel Partial Emulation (VPPE) that utilizes the Scaled Vecchia approximation within the PPE framework to allow for parallel partial emulation with larger training datasets. The VPPE is applied to three computer experiments, a synthetic data set, a hydrology model, and a volcanic flow model, yielding comparable predictive accuracy to the PPE at a fraction of the runtime.

Joshua Seidman
Marquette University
josh.seidman@marquette.edu

MS162

Developing a Field-Scale Digital Twin for the Land-Use Response to Agricultural and Environmental Policies in the UK and Giving It to Policymakers

Understanding the impacts of land-use change in a single field requires well calibrated predictive models for the existing land-use and the potential. Policymakers must hit statutory environmental targets in carbon sequestration, reversing biodiversity loss, water quality and food/energy security. Their use of the land to achieve these goals is limited by the fact that they dont own the land and can only offer subsidies to the few hundred thousand farmers that do. Our group tries to offer effective decision support to policymakers, model all the current and possible future uses of all fields in all farms in the UK at the same time, quantify uncertainty in the response of the farmers and hence the landscape to any particular policy, and wrap all of that in tools fast enough to enable policymakers to explore their options. Is that even possible? In this talk we look at our attempts to deliverthis kind of UQ to the UK government in the last 12 months.

Daniel Williamson
University of Exeter
d.williamson@exeter.ac.uk

MS163

A Unified Probabilistic Programming Framework for State Estimation and Dynamics Discovery

Many critical forecasting problems in science and engineering involve partially observed systems with unknown underlying dynamics. Despite their importance, current methodologies for learning such systems remain fragmented across specialized, often incompatible computational frameworks and inference paradigms. This talk introduces a novel probabilistic programming framework for sequential inference that provides a unifying modeling and inference abstraction, enabling joint Bayesian estimation of latent states and governing dynamical laws. By integrating state-space models, stochastic differential equations (SDEs), and machine learning components with filter- and smoother-based approximate likelihoodsincluding ensemble Kalman and particle filtersthe framework facilitates a rigorous approach to system identification. We demonstrate the utility of this interface across SDEs, hidden Markov models, and general discrete-time transition systems, showing its capacity to learn dynamics represented by neural networks, mechanistic equations, sparse functional libraries, and hybrid symbolic-neural architectures. Finally, we illustrate how these tools are wrapped in principled Bayesian uncertainty quantification and model comparison to ensure robust, auditable dynamical forecasting.

Matt Levine
Basis
matt@basis.ai

Iñigo Urteaga
Basque Center for Applied Math
Bilbao, Spain
iurteaga@bcamath.org

Dan Waxman
Basis

dan@basis.ai

MS163

PyDPF: A Python Package for Differentiable Particle Filters

State-space models (SSMs) are a widely used tool in time series analysis. In the complex systems that arise from real-world data, it is common to employ particle filtering (PF), an efficient Monte Carlo method for estimating the hidden state corresponding to a sequence of observations. Applying particle filtering requires specifying both the parametric form and the parameters of the system, which are often unknown and must be estimated. Gradient-based optimisation techniques cannot be applied directly to standard particle filters, as the filters themselves are not differentiable. However, several recently proposed methods modify the resampling step to make particle filtering differentiable. In this talk, we present an implementation of several such differentiable particle filters (DPFs) with a unified API built on the popular PyTorch framework. Our implementation makes these algorithms easily accessible to a broader research community and facilitates straightforward comparison between them. We validate our framework by reproducing experiments from several existing studies and demonstrate how DPFs can be applied to address several common challenges with state space modelling.

John-Joseph Brady
King's College London
john-joseph.brady@kcl.ac.uk

Benjamin Cox
University of Edinburgh
benjamin.cox@ed.ac.uk

Yunpeng Li
King's College London
yunpeng.li@kcl.ac.uk

Víctor Elvira
University of Edinburgh
victor.elvira@ed.ac.uk

MS163

GP4SP: Online and Sequential Inference of Gaussian Processes for Signal Processing

The proliferation of capable and efficient machine learning (ML) models marks one of the strongest methodological shocks to signal processing (SP) in its nearly 100-year history. ML models allow complex, nonlinear, and accurate SP systems, yet a common hurdle in adapting them is the necessity of sequential inference, which differs both theoretically and methodologically from the usual train/test-split batch paradigm. Gaussian processes (GPs) are a flexible yet principled framework to model random functions, whose influence has naturally extended to SP as statistical and ML methods become increasingly central. We provide a self-contained, tutorial-style overview of GPs, with particular focus on recent methodological advances in the sequential, incremental, or streaming inference of GPs, and on sequential inference using GPs. We introduce these techniques from a signal-processing perspective while bridging to recent advances in ML. Many of the developments we survey have direct applications to state-space modeling, sequential regression and forecasting, anomaly detection in time series, sequential Bayesian optimization, adaptive and active sensing, and sequential detection and decision-making. By unifying these advances under a signal-processing lens, we aim to equip practitioners with practical tools and a coherent roadmap for deploying sequential GP models in real-world systems.

Daniel Waxman
Stony Brook University
dan@basis.ai

Fernando Llorente
Brookhaven National Laboratory
fllorente@bnl.gov

Petar Djuric
Stony Brook University
petar.djuric@stonybrook.edu

MS164

Stochastic Data-Driven Reduced Order Models for Turbulent Flows

In this talk, I will survey our group's work on the development of stochastic data-driven reduced order models (ROMs) for turbulent flows. In particular, I will discuss the construction of ROM closures and stabilizations that leverage data-driven modeling, spatial filtering, stochastic modeling, and data assimilation. The new ROMs will be assessed in the numerical simulation of challenging test problems that are centered around chaotic fluid flows, such as the turbulent channel flow and the quasi-geostrophic equations that model the large scale ocean circulation.

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MT1

Diffusion Models for Inverse Problems: Theory and Application

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Assad A. Oberai
University of Southern California
aoberai@usc.edu

MT2

Conformal Prediction - An Introductory Tutorial and Recent Advances

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Aymeric Dieuleveu
École Polytechnique, France
aymeric.dieuleveut@polytechnique.edu

MT3

Causal Reasoning in Generative Models for Artifi-

cial Intelligence

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Bryon Aragam
University of Chicago
Nikhyl.Aragam@chicagobooth.edu

MT4

Mathematical and Statistical Foundations of Global Sensitivity and Probabilistic Robustness Analyses

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Fabrice Gamboa
Institut de Mathématiques de Toulouse. AOC Project
University of Toulouse
fabrice.gamboa@math.univ-toulouse.fr

PP1

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

We study iterative solvers for unsteady incompressible NavierStokes flow with uncertain viscosity. We formulate an all-at-once stochastic Galerkin system, where the systems across all time steps are solved simultaneously. Nonlinearity is handled via Picard linearization, and low-rank tensor-train representations of solutions are employed to reduce memory and computational costs. To control the rapid rank growth induced by the convection operator, we approximate its CANDECOMP/PARAFAC (CP) decomposition using an alternating least-squares (ALS) procedure. This enables the construction of structure-preserving preconditioners that more effectively capture both stochastic and nonlinear dynamics compared to traditional mean-based approaches. The effectiveness of the method is illustrated by numerical experiments.

Ahmet Kaan Aydin
University of Maryland - Baltimore County
aaydin@umbc.edu

Bedrich Sousedik
University of Maryland Baltimore County
sousedik@umbc.edu

PP1

Local Dynamic Calibration Via Just-in-Time Split Conformal Prediction in Latent Spaces

We address the challenge of applying conformal prediction (CP) to online soft sensing of product quality in batch processes, where the key assumption of data exchangeability is violated due to temporal dependence and evolving operating conditions. We propose a mechanism that integrates Split Conformal Prediction (SCP) with Just-in-Time Learning (JITL) in a latent space constructed via self-organizing maps (SOM). At each prediction time, a buffer of past latent representations is collected based on (i) topological similarity and (ii) distance-based similarity. A new model is then trained from scratch on a randomly partitioned buffer: one subset for model fitting and one for calibration. SCP quantiles are derived over multiple splits, yielding distribution-free prediction intervals tailored to the current context. This procedure is repeated for every batch, resulting in a dynamic, locally calibrated prediction mechanism. We investigate whether local exchangeability can be formalized in latent spaces with such a selection mechanism and explore conditions under which theoretical coverage guarantees may be derived. The approach is evaluated on a Michelin rubber mixing process, highlighting its potential to provide reliable uncertainty quantification in non-exchangeable, time-dependent industrial systems.

Louis Berthier
Michelin
Centre de Mathématiques Appliquées
louis-desire-romeo.berthier@michelin.com

Maxime Moreaud, Guillaume Ramelet
Michelin
maxime.moreaud@michelin.com,
guillaume.ramelet@michelin.com

Ahmed Shokry
Centre de Mathématiques Appliquées, Ecole Polytechnique
ahmed.shokry@polytechnique.edu

Aymeric Dieuleveut
Ecole Polytechnique, France
aymeric.dieuleveut@polytechnique.edu

PP1

Bayesian Nonparametric Calibration of Climate Models

A climate model (CM) is a set of differential equations representing the behavior of the Earth system. For a given input-parameter value, the CM encodes a distribution that is analytically intractable but samples of which can be produced at high computational cost by running the CM with slightly perturbed initial conditions. The goal is to calibrate the CM by finding the parameter value that most closely matches the CM to the true climate distribution. Existing methodologies achieve this calibration task by fitting Gaussian process to the CM samples that only targets the mean and does not address the distributional uncertainty present within the climate model. Here, we propose a calibration algorithm via Bayesian optimization using a flexible probabilistic surrogate model that can learn any arbitrary distribution through a random bijective mapping between the CM samples and standard normal distribution. We further model the bijective maps using autoregressive Gaussian processes whose inputs include the CM input parameters. Our framework is scalable and hence it can be applied to large spatial climate fields. We demonstrate the performance of the approach on numerical examples.

Anirban Chakraborty
Medical University of South Carolina
chakraan@musc.edu

Paul Wiemann
The Ohio State University
wiemann.3@osu.edu

Daniel Drennan
Department of Statistics, Texas A&M University

danjdrennan@tamu.edu

Felix Jiminez
Amazon
felix.m.jimenez3@gmail.com

Matthias Katzfuss
University of WisconsinMadison
katzfuss@gmail.com

PP1

Projected Score Matching for Generative Modeling and Inverse Problems in Chaotic Systems

Physical measures associated with chaotic systems are often singular but admit a conditional density on the lower-dimensional unstable manifold. In this work, we exploit this lower-dimensional absolutely continuous structure to enable faster generative modeling of such measures. In particular, the learned drift term or vector field of the generative model only uses the well-defined unstable component of the score. To supply the training data for the unstable scores, we use an existing, rigorous algorithm from Chandramoorthy and Wang SIAM J. Appl. Dyn. Sys. 2022. Using these conditional scores, and numerically estimated unstable tangent vectors, we use the learned vector field and a reconstruction scheme to sample from the target measures. We show preliminary numerical results on low-dimensional chaotic systems, wherein we obtain improved reconstructions compared to a vanilla score-based model for the same number of training examples.

Sebastian Deery, Nisha Chandramoorthy
The University of Chicago
sbdeery@uchicago.edu, nishac@uchicago.edu

PP1

Phase Reduction in Particle Method for Simulation of Stochastic Limit Cycle Oscillators

Noisy oscillators arise in the modeling of many scientific problems, particularly in physics, neuroscience, ecology and climate studies. Phase reduction is a traditional approach that simplifies the dynamics. However, the assumption of weak noise in phase reduction limits its application. Particle methods, on the other hand, provide more accurate results, but are often computationally expensive. In this work, we present how phase reduction techniques can be incorporated into particle methods to study population level behaviors of noisy coupled oscillators. We use elliptical Gaussian basis functions to solve the Fokker-Planck equation that describes the evolution of the probability density function. Previous studies have suggested splitting and combining the Gaussian kernels under certain conditions. We propose an original particle shrinking method based on phase reduction. Our method reduces the number of particles and accelerates the particle updating process, thereby achieving higher computational efficiency. At the same time, we preserve the structure of the possibly high-dimensional state space, which avoids over-simplification. In particular, we show that this method is well adapted to systems with relaxation oscillation. As illustration, we apply our method to 4-dimensional and 5-dimensional Hodgkin-Huxley type neuronal models.

Yangyang Du, Daniel Forger
University of Michigan
yyadu@umich.edu, forger@umich.edu

PP1

SIPF-PIC: Efficient Large-Scale Simulation of 3D Chemotactic Systems

The Keller-Segel system models biological aggregation driven by chemotaxis, such as bacterial pattern formation. Simulating its complex 3D dynamics, especially near blow-up singularities where density diverges, is computationally prohibitive with existing methods. To overcome this, we developed the SIPF-PIC framework, a novel hybrid method combining Lagrangian particle tracking with spectral field solvers using optimized interpolation. Crucially, SIPF-PIC reduces the cost per timestep from $\mathcal{O}(PH^3)$ to $\mathcal{O}(P + H^3 \log H)$ while maintaining convergence, where P is the particle count and H the grid resolution per dimension. This enables large-scale simulations capturing elusive blow-up structures like evolving ring singularities, beyond simple point collapse.

Jingyuan Hu
The University of Hong Kong
Abbhl126@gmail.com

PP1

Total HSIC: A New Index for Global Sensitivity Analysis

Global sensitivity analysis (GSA) is a fundamental component of uncertainty quantification used to identify which input parameters most influence a model's output. Traditional methods, such as variance-based Sobol' indices, are limited because they only capture influence that manifests through variance and assume statistical independence of inputs. This can lead to inaccurate sensitivity scores, particularly when dealing with complex, non-linear dependencies. We propose an alternative approach using the Hilbert-Schmidt Independence Criterion (HSIC). Unlike variance-based methods, HSIC is a non-parametric measure that captures general statistical dependence, including non-linear and higher-order effects, by embedding probability distributions into Reproducing Kernel Hilbert Spaces (RKHS). Our findings demonstrate that this approach offers a compelling and flexible alternative to traditional GSA methods, capable of detecting a wider range of statistical dependencies.

Troy Larsen
North Carolina State University
tjlarse2@ncsu.edu

Alen Alexanderian
North Carolina State University, U.S.
alexanderian@ncsu.edu

PP1

Regularity Estimate and Sparse Approximation of Pathwise Robust Duncan-Mortensen-Zakai Equation

In this paper, we establish *a priori* estimate for arbitrary-order derivatives of solution to pathwise robust Duncan-Mortensen-Zakai equation within the framework of weighted Sobolev space. The weight function, which vanishes on the physical boundary, is crucial in completing the *a priori* estimate proof, but also introduce a loss of regularity estimate near the boundary. Therefore, we em-

ploy the Sobolev inequalities and their weighted analogues to sharpen the regularity bound, providing improvements in both classical Sobolev space and Hölder continuity estimates. The refined regularity estimate reinforces the plausibility of the quantized tensor train (QTT) method and consolidates the conditions under which the method converges. To further enhance the capacity of the method to solve nonlinear filtering problem in real time manner, we refine the method in the case of a functional polyadic drift f and observation h. Finally, we perform some numerical simulations to reaffirm our theory. For high-dimensional cubic sensor problems, our method demonstrates superior efficiency and accuracy in comparison to the particle filter (PF) and extended Kalman filter (EKF). Beyond this, for multi-mode problems, while PF exhibits lack of precision due to its stochastic nature and EKF is constrained by the Gaussian assumption, the enhanced method provides exact reconstruction of multi-mode conditional density function.

Yuhua Meng
The University of Hong Kong
yvette00@connect.hku.hk

Zhongjian Wang
Nanyang Technological University
zhongjian.wang@ntu.edu.sg

Stephen S.T. Yau
Tsinghua University
yau@uic.edu

Zhiwen Zhang
Department of Mathematics,
The University of Hong Kong,
zhangzw@hku.hk

PP1

Uncertainty Quantification in Tracer Transport: Pressure-to-concentration Propagation vs. pCN-DE McMC Inference

Uncertainty quantification in porous media flow remains a major challenge due to the high dimensionality of the parameter space, which strongly limits the efficiency of standard Markov chain Monte Carlo (McMC) methods. Even after dimensionality reduction, achieving robust sampling with fast convergence is far from trivial, and the difficulty becomes even more pronounced in transient problems. In previous work, the authors showed that the Preconditioned CrankNicolson Differential Evolution (pCN-DE) algorithm yields satisfactory performance for 2D single-phase flow problems with five-spot well configurations. Building upon these results, we now investigate uncertainty propagation for solute transport by considering concentration fields in a 2D tracer problem. Specifically, we compare two approaches: (i) propagating uncertainty from the pressure solutions to the concentration, and (ii) performing a direct McMC simulation on the concentration data using the pCN-DE method. The comparison highlights the differences between propagated and directly inferred statistics, providing insights into the efficiency and accuracy of each approach.

Michel Tosin, Marcio Borges
LNCC/MCTI
michelt@posgrad.lncc.br, mrborges@lncc.br

Felipe Pereira
The University of Texas at Dallas
luisfelipe.pereira@utdallas.edu

PP1

Numerical Approximation of Effective Diffusivities in Homogenization of Nondivergence-Form Equations with Large Drift by a Lagrangian Method

In this paper, we study numerical methods for the homogenization of linear second-order elliptic equations in nondivergence-form with periodic diffusion coefficients and large drift terms. Upon noting that the effective diffusion matrix can be characterized through the long-time variance of an associated diffusion process, we construct a Lagrangian numerical scheme based on a direct simulation of the underlying stochastic differential equation and utilizing the framework of modified equations, thereby avoiding the need to solve the Fokker–Planck–Kolmogorov equation. Through modified equation analysis, we derive higher-order weak convergence rates for our method. Finally, we conduct numerical experiments to demonstrate the accuracy of the proposed method. The results show that the method efficiently computes effective diffusivities, even in high dimensions.

Han Wu
University of Hong Kong
shivers641@gmail.com

PP1

Activation Level Uncertainty Quantification of Hyperspectral Neural Radiance Fields for High Consequence Remote Sensing Applications

Hyperspectral imaging (HSI) is a key capability for detection of targets such as gases, liquids, or solids in remote sensing applications. The detection process becomes challenging when objects in the scene occlude the target. We propose to advance the utility of 3D scene understanding methods, neural radiance fields (NeRFs), to cohesively synthesize multi-observation HSI data into a single scene thereby providing the ability to look around targets or obstructions. NeRFs have shown promising results in addressing this task, specifically for gas plume detection and identification, but the limitations of these newer models have not been fully explored. We leverage Gaussian processes to capture neuron level uncertainty per layer. By quantifying uncertainty at the the activation level, we gain insight into the stability of the learned representation at each stage of the model. The representation level uncertainty highlights where the network is confident in its spectral-spacial encoding versus where it may be underconstrained. For remote sensing applications, decision-makers require detection outputs with measures of confidence in those outputs, especially in complex environments where occlusions, atmospheric effects, or sparse measurements can degrade performance. Understanding and propagating representation uncertainty allows us to flag unreliable predictions in a single off-the-shelf model, prioritize additional observations, and ultimately improve trust and robustness.

Sophia Zorek
Rice University
saz2@rice.edu

Zig Hampel-Arias, Adra Carr
Los Alamos National Laboratory

zhampel@lanl.gov, acarr@lanl.gov