Mathematical Foundations & Methodologies for Artificial Intelligence & Data-Driven Scientific Computing
Workshop on Data-Drive Scientific Computing
Abstracts

11 Aug 2026–15 Aug 2026

July 22, 2026

TBA

15 Aug
9.30 am

Zhi Chen
Peking University, China

Pending

Foundation Models for Physics: From Transformers to Kinetic and Self-Warping Architectures

11 Aug
1.30 pm

Maarten de Hoop
Rice University, USA

We present a kinetic theory perspective of foundation models for physics. We begin with providing a mathematical framework for analyzing transformers. To uniformly address their expressivity, we consider the case that the mappings are conditioned on a context represented by a probability distribution of tokens. That is, transformers become mappings between probability measures. We demonstrate that deep transformers are universal and can approximate continuous in-context mappings to arbitrary precision, uniformly over compact token domains. We then characterize the conditions on mappings between measures that enable these to be represented in terms of in-context mappings as transformers through push forwards. The solution map of the Vlasov equation for interacting particle systems in the mean-field regime for the Cauchy problem satisfies the conditions; conversely, we prove that the measure-theoretic self-attention has the properties that ensure that the infinite depth, mean-field transformer can be identified with a Vlasov flow. Extending this framework from interactions to collisions leads to a further development of structured architectures inspired by Lattice Boltzmann Models. A complementary direction is provided by an architecture built from multihead self-warping operators. Each head predicts, using local state information, a displacement field and pulls back features from an adaptively selected upstream location.

Joint research with I. Dokmanić, T. Furuya, K. Hegazy, A. Lara Benitez, M. Lassas, M. Mahoney, T. Muser, G. Peyré, A. Spitzer and T. Teolis

---

Deep Learning for Reynolds-Averaged Navier-Stokes Simulations of Turbulent Flows

12 Aug
9.30 am

Daniel Dehtyriov

University of Oxford, UK

Turbulence is ubiquitous in engineering and science, yet direct simulation is prohibitively expensive. The Reynolds-averaged Navier–Stokes (RANS) equations are far cheaper but introduce unclosed terms (the closure problem). Offline-trained machine-learning closures suffer distribution shift in predictive simulations, while methods that bypass the governing equations struggle to generalise from scarce data. We develop a physics-derived deep learning closure for RANS, the Deep Algebraic Reynolds Stress Model (DARSM), in which a neural network sets the empirical parameters of an implicit algebraic Reynolds stress equation, so it learns only a small physics-constrained correction. Training end-to-end through the coupled solver removes distribution shift but defeats standard automatic differentiation; we derive adjoint equations that reuse the solver's implicit solves in transpose, a construction that extends to any semi-implicit PDE solver with embedded neural networks. DARSM cuts out-of-sample velocity error 2-4x over baseline RANS and transfers without retraining from attached (square duct) to separated (periodic hills) flow, outperforming five established ML methods. The broader lesson is a template for scientific machine learning: constrain the network with the physical structure the equations already provide, and train it through the solver rather than against precomputed data.

Flow-based Bayesian Filtering for High-dimensional Nonlinear Stochastic Dynamical Systems

15 Aug
2.30 pm

Ling Guo

Shanghai Normal University, China

Bayesian filtering for high-dimensional nonlinear stochastic dynamical systems is a fundamental yet challenging problem in many fields of science and engineering. Existing methods face significant obstacles: Gaussian-based filters struggle with non-Gaussian distributions, while sequential Monte Carlo methods are computationally intensive and prone to particle degeneracy in high dimensions. Although generative models in machine learning have made significant progress in modeling high-dimensional non-Gaussian distributions, their inefficiency in online updating limits their applicability to filtering problems. In this talk, we will present a flow-based Bayesian filter (FBF) that integrates normalizing flows to construct a novel latent linear state-space model with Gaussian filtering distributions. This framework facilitates efficient density estimation and sampling using invertible transformations provided by normalizing flows, and it enables the construction of filters in a data-driven manner, without requiring prior knowledge of system dynamics or observation models. Numerical experiments demonstrate the superior accuracy and efficiency of the FBF method.

Curvature-Aware Optimization for High-Accuracy Physics-Informed Neural Networks

13 Aug
2.30 pm

George Em Karniadakis

Brown University, USA

Efficient and robust optimization is essential for neural networks, enabling scientific machine learning models to converge rapidly to very high accuracy – faithfully capturing complex physical behavior governed by differential equations. In this work, we present advanced optimization strategies to accelerate the convergence of physics-informed neural networks (PINNs) for challenging partial (PDEs) and ordinary differential equations (ODEs). Specifically, we provide efficient implementations of the Natural Gradient (NG) optimizer, Self-Scaling BFGS and Broyden optimizers, and demonstrate their performance on problems including the Helmholtz equation, Stokes flow, inviscid Burgers equation, Euler equations for high-speed flows, and stiff ODEs arising in pharmacokinetics and pharmacodynamics. Beyond optimizer development, we also propose new PINN-based methods for solving the inviscid Burgers and Euler equations, and compare the resulting solutions against high-order numerical methods to provide a rigorous and fair assessment. Finally, we address the challenge of scaling these quasi-Newton optimizers for batched training, enabling efficient and scalable solutions for large data-driven problems.

Accelerated Langevin Sampling of Complex Distribution using Variable Timestep and Friction

14 Aug
9.30 am

Benedict Leimkuhler

University of Edinburgh, UK

I will describe a new framework for sampling high-dimensional and multi-modal probability distributions using a Langevin scheme that combines SamAdams variable stepsize with a variable damping matrix that concentrates friction in the direction of the local force. By careful construction the scheme requires essentially no additional overhead compared to standard Langevin methods and it can be shown to substantially accelerate convergence to the target distribution with no compromise in accuracy. This is joint work with Peter A. Whalley (ETH Zurich).

Learning Mesoscopic Dynamics

12 Aug
2.30 pm

Qianxiao Li
National University of Singapore, Singapore

We discuss some recent work on constructing stable and interpretable mesoscopic dynamics from trajectory data using deep learning. We adopt a hypothesis-driven approach: contrary to using generic neural networks as functional approximators for the vector fields driving a dynamical system or operator based learning approaches, we start with a structured hypothesis class in which well-posedness and numerical stability are a priori ensured. Learning on trajectory data then identifies a specific member of this family suitable for capturing the observed dynamical process, from which scientific predictions and analysis can be made. We demonstrate this approach on a variety of mesoscopic modelling problems, including the dynamics of spin systems and self-healing in high-entropy alloys.

Foundation Neural Operators and Diffusion Models over Function Spaces

11 Aug
9.30 am

Lu Lu
Yale University, USA

Deep neural operators have emerged as a powerful paradigm in scientific machine learning (SciML) for learning nonlinear mappings between function spaces arising in complex dynamical systems. In this talk, I will introduce the deep operator network (DeepONet) and its extensions, such as geometry-dependent/manifold operator learning, and demonstrate their effectiveness on diverse multiphysics and multiscale 3D problems, such as geological carbon sequestration, fire simulation, full waveform inversion, and topology optimization. I will then present the first operator learning method requiring only a single PDE solution, i.e., one-shot learning, enabled by a new concept of local solution operator motivated by the locality of physics. Moreover, I will discuss our recent work on diffusion models, including FunDiff as a novel framework of diffusion models over function spaces for physics-informed generative modeling, and RED-DiffEq which leverages denoising diffusion models as regularization for inverse PDE problems. Finally, I will present a physics-informed multimodal foundation model for PDEs and neural-operator element method (an efficient and scalable finite element method enabled by reusable neural operators).

In-Context Learning in Scientific Computing

11 Aug
2.30 pm

Yulong Lu
University of Minnesota, USA

Transformer-based foundation models, pre-trained on large datasets spanning a wide range of tasks, have shown remarkable adaptability to diverse downstream applications—even in low-data regimes. A particularly striking capability is in-context learning (ICL): when given a prompt containing a few examples from a new task alongside a query, these models can produce accurate predictions without any parameter updates. This emergent behavior is often viewed as a paradigm shift for transformers, yet its theoretical foundations remain only partially understood. In this talk, I will present recent theoretical progress toward understanding ICL in scientific computing. I will focus on understanding how transformer architectures can implicitly perform task adaptation in three representative problem classes: learning solution operators of PDEs, dynamical system prediction and generative modeling.

Efficient Sampling, Optimization, and Targeted Diffusion Generation for Physical AI

11 Aug
10.45 am

Yueming Lyu
*A*STAR, Singapore*

This talk presents a unified perspective on efficient sampling, black-box optimization, neural operators, and targeted generation for Physical AI. The central theme is the development of data-efficient and computation-efficient methods for exploring complex spaces, learning functional mappings, and generating solutions that satisfy desired objectives or constraints.

The first part of the talk discusses quasi-Monte Carlo sampling as a structured alternative to random sampling, with applications in numerical integration, uncertainty quantification, and simulation-based learning. In particular, a novel subgroup rank-1 lattice sampling method will be introduced to improve sampling efficiency and coverage. The talk then connects efficient sampling with black-box optimization, where the goal is to identify high-quality solutions when gradients, explicit models, or closed-form objective functions are unavailable. Along this direction, a fast rank-1 lattice targeted sampling method will be presented for query-efficient black-box optimization.

Furthermore, the talk introduces an infinite-order kernel neural operator for learning solution operators of partial differential equations, with the goal of improving the expressivity and generalization of neural operator models in physical systems.

Finally, guided-diffusion targeted generation will be discussed for design and discovery, where generative models are steered toward producing samples with specified physical properties. Overall, the talk highlights how efficient sampling, optimization, operator learning, and guided generation can jointly support scalable and controllable Physical AI.

Flow-based Generative Models for Amortized Bayesian Inference in Regression and Inverse PDE Problems

12 Aug
4.00 pm

Xuhui Meng

Huazhong University of Science and Technology, China

Bayesian inference provides a principled framework for uncertainty quantification in scientific machine learning. However, conventional Bayesian approaches typically require solving a new inference problem for each observation set, resulting in substantial computational costs that hinder real-time applications such as online monitoring, digital twins, and decision-making under uncertainty. Furthermore, many scientific machine learning problems involve inference over infinite-dimensional function spaces with varying numbers and locations of observations, posing significant challenges for existing amortized inference methods. In this work, we propose Flow-ABI, a flow-based generative framework for amortized Bayesian inference in regression and inverse partial differential equation (PDE) problems. The proposed framework consists of two components: (i) a functional prior model that learns expressive prior in function spaces directly from historical data and physical knowledge through flow matching, and (ii) a set-conditioned functional posterior sampler that performs amortized Bayesian inference by mapping observation sets to posterior distributions over functions. The learned posterior model naturally accommodates varying numbers and locations of observations, is invariant to permutations of the measurement set, and generalizes across different observation discretizations. Once trained, Flow-ABI enables near-real-time posterior sampling for previously unseen observations without retraining or iterative optimization. The proposed methodology can be seamlessly integrated with a broad class of scientific machine learning frameworks, including physics-informed neural networks and neural operators, for uncertainty-aware inverse PDE modeling. Numerical experiments on representative regression and inverse PDE benchmarks demonstrate that Flow-ABI accurately captures both Gaussian and non-Gaussian posterior distributions while achieving speedups of more than two orders of magnitude relative to the gold-standard Bayesian inference method, Hamiltonian Monte Carlo. These results demonstrate that the Flow-ABI as an effective, scalable, and computationally efficient framework for uncertainty quantification in scientific machine learning.

Intelligent Scientific Computing: from Data-driven Modeling to Discovery

15 Aug
10.45 am

Hao Sun

Renmin University, China

Scientific computing plays a crucial rule in characterizing the behavior of complex systems in many disciplines. Although AI techniques, particularly deep learning, demonstrate significant potential, they suffer from critical limitations including limited generalizability, uncontrollable errors, and strong dependency on data. Consequently, it is imperative to integrate data-driven machine learning with knowledge-driven symbolic computation to advance novel intelligent scientific computing theories and methodologies. This talk will discuss the challenges of modeling and simulating complex spatiotemporal dynamics, such as turbulence, climate systems, and reaction-diffusion processes. This talk will also introduce a series of physics-inspired and symbolic learning methods for complex system modeling, accelerated simulation, optimal design, and knowledge discovery.

Kolmogorov Superpositions: Construction, Approximation and Networks

14 Aug
10.45 am

Li-Lian Wang

Nanyang Technological University, Singapore

The Kolmogorov Superposition Theorem (KST, 1957), offers a mathematically elegant framework for expressing any high-dimensional continuous function as a superposition of lower- or one-dimensional continuous functions. This foundational theorem has recently gained renewed interest, particularly in neural networks. However, a major challenge remains: the one-dimensional functions resulted from all constructions are highly non-smooth. In this talk, we present a novel approximate version of KST involving C^2 inner functions and piecewise C^2 outer functions and show its applications in neural network approximations.

This talk will be based on joint works with Z. Chen (NTU), L. Song (Lanzhou University) and J. Toscane and G. Karniadakis (Brown University).

Preconditioned One-step Generative Modelling for Bayesian Inverse Problems

12 Aug
1.30 pm

Zhongjian Wang

Nanyang Technological University, Singapore

We propose a machine-learning algorithm for Bayesian inverse problems in the function-space regime based on one-step generative transport. Building on the Mean Flows, we learn a fully conditional amortized sampler with a neural-operator backbone that maps a reference Gaussian noise to approximate posterior samples. We show that while white-noise references may be admissible at fixed discretization, they become incompatible with the function-space limit, leading to instability in inference for Bayesian problems arising from PDEs. To address this issue, we adopt a prior-aligned anisotropic Gaussian reference distribution and establish the Lipschitz regularity of the resulting transport. Our method is not distilled from MCMC: training relies only on prior samples and simulated partial and noisy observations. Once trained, it generates a 64×64 posterior sample in $\sim$ 1e$-$3s, avoiding the repeated PDE solves of MCMC while matching key posterior summaries.

TBA

12 Aug
10.45 am

Yang Xiang

Hong Kong University of Science and Technology, Hong Kong SAR

Pending

A Hybrid Two-level MCMC Framework to Accelerate Posterior Mean Estimation with Deep Learning Surrogates for Bayesian Inverse Problems

15 Aug
1.30 pm

Juntao Yang

NVIDIA, Singapore

Bayesian inverse problems arise in various scientific and engineering domains, and solving them can be computationally demanding. This is especially the case for problems governed by partial differential equations, where the repeated evaluation of the forward operator is extremely expensive. Recent advances in Deep Learning (DL)-based surrogate models have shown promising potential to accelerate the solution of such problems. However, despite their ability to learn from complex data, DL-based surrogate models sometimes may not match the accuracy of high-fidelity numerical

models, which limits their practical applicability. We propose a hybrid two-level Markov Chain Monte Carlo (MCMC) method that combines the strengths of DL-based surrogate models and high-fidelity numerical solvers to compute the posterior mean of Quantities of Interest (QoI) in Bayesian inverse problems governed by partial differential equations. The intuition is to leverage the evaluation speed of a DL-based surrogate model as the base chain, and correct its errors using a limited number of high-fidelity numerical model evaluations in a correction chain, hence the name hybrid two level MCMC. The proposed hybrid framework can be generalized to other approaches such as the ensemble Kalman filter and sequential Monte Carlo methods.

TBA

13 Aug
1.30 pm

Liu Yang
National University of Singapore, Singapore

Pending

When Does Scientific AI Actually Use Physics? From Causal Mechanism Audits to Executable Solver Semantics

Xinyu Yang
National University of Singapore, Singapore

Data-driven scientific computing increasingly incorporates physical structure through inductive biases, auxiliary variables, differentiable solvers, and tool-using agents. Yet the presence of such structure does not guarantee that a model genuinely relies on it, nor does access to simulation software imply an understanding of the numerical methods being executed.

This talk brings together two complementary studies that address these limitations. The first develops a framework for testing whether machine learning models causally use the mechanisms provided to them. Using structured auxiliary interfaces and controlled interventions, the analysis separates genuine dependence from superficial references and shortcut-based prediction. The results show that improved accuracy, interpretable intermediate outputs, or explicit mention of a mechanism do not, by themselves, demonstrate that the mechanism influenced the model's reasoning. The second study introduces LOCUS, a cross-discretization semantic contract for representing trusted local PDE operations as typed, validated, and differentiable computational objects. Its gather-local-scatter abstraction captures FDTD stencils, FEM residuals, and FVM fluxes while preserving the semantics of the underlying discretization. This representation enables gradient auditing, multiphysics composition, backend-independent execution, structure-preserving learned closures, and validator-gated modifications by AI agents.

Taken together, the two studies suggest a broader direction for trustworthy scientific AI. Physical structure should be accompanied by evidence of causal use, while simulation systems should expose numerical semantics that can be inspected, composed, differentiated, and safely modified. This combination provides a foundation for more reliable data-driven solvers and bounded agents for scientific computing.

TBA

13 Aug
9.30 am

Pin Zhang

National University of Singapore, Singapore

Pending

TBA

13 Aug
10.45 am

Hongkai Zhao

Duke University, USA

Pending
