Workshop on Learning and Computational Science
Abstracts

13 Mar 2026–15 Mar 2026

March 6, 2026

A Systems-Level Environment for Scientific Computing, Coding and Visualization

David Adalsteinsson

University of North Carolina at Chapel Hill, USA

ImageTank is a macOS-based scientific computing environment that combines high-level visual programming with systems-level design including automatic memory management, multi-threaded processing, file sharing, and multi-core parallelism.

Unlike conventional graphing or image processing applications as well as web environments like Jupyter notebooks, ImageTank operates as a computational environment where users can build and iterate on analysis pipelines without writing traditional code. When deeper customization is needed, ImageTank can be extended to external programming languages including C++, Python, and Julia, and the integration glue required to connect them is automatically handled by the system. This allows researchers to incorporate existing code or specialized libraries where needed.

This presentation draws on several published collaborations at the University of North Carolina to illustrate ImageTank's application across different research contexts: processing lattice light-sheet microscopy images of FRET biosensors, automated detection and analysis of exocytosis events, and polymer simulation of chromosomal DNA dynamics interfaced with C++ code. Together these examples reflect the range of applications that the environment has been applied to, from fluorescence microscopy to computational modeling.

Attendees will come away with an understanding of ImageTank's core architecture, practical examples of how it has supported real research workflows, and a sense of how it might apply to their own data challenges.

On the Accuracy of Tensor Transfer Method for Open Quantum Systems

Zhenning Cai

National University of Singapore, Singapore

The dynamics of open quantum systems is governed by the Nakajima-Zwanzig (NZ) equation, an integral equation in which the non-Markovian effect is described by a convolutional term. In general, the convolutional kernel is unknown, and therefore one solution is to first compute the solution using path integrals, and then "learn" the kernel based on a discrete form of the NZ equation. This approach is known as the tensor transfer method. However, recently, there is a series of debate on the validity of the tensor transfer method about its accuracy. Two research groups are holding conflicting opinions on whether a desired order of accuracy can be achieved using TTM. To better understand the issue, we started with the diagrammatic representation of open quantum system dynamics, and discovered a TTM-like structure for the exact solution, allowing us to the compute the discrete kernel directly without the learning procedure. The desired numerical order with the TTM form is therefore guaranteed.

Moving Artificial Intelligence Algorithms for Pathology Closer to the Clinic by Addressing Technical Barriers to Scale

Iain Carmichael

University of North Carolina at Chapel Hill, USA

Artificial intelligence (AI) algorithms show significant potential for histopathology imaging data (high-resolution images of diseased tissue studied under a microscope). Systems leveraging AI algorithms may increase access to state-of-the-art care, improve diagnostic precision, enhance clinical decision making, and advance the basic scientific investigation of disease processes. Several key challenges currently hold back our ability to leverage AI algorithms to improve patient care including: the cost of obtaining expert image annotations to train segmentation algorithms, the complexity / sheer size of Pathology imaging (e.g. one image can be up to 100,000x100,000 pixels in dimension), and the fact that much of the useful information contained in electronic medical records is stored in free-text notes. We discuss several approaches we are weaving together — including software systems and novel algorithms — to address these key challenges.

Preparing for Catastrophe in a Changing World: The Role of Data Science and Machine Learning in the (re)Insurance Sector

Zeff Harrison

University of North Carolina at Chapel Hill, USA

The global (re)insurance industry manages risks at an almost unimaginable scale, with policy premiums representing over 10% of GDP in several countries across Europe, Asia, and North America. Emerging perils tied to climate change, demographic patterns, and advances in information technology make it more difficult for insurers to use traditional actuarial methods to quantify risk and set prices. This uncertainty is causing insurers to stop covering some kinds of risk altogether, destabilizing markets and making it more difficult for communities and critical industries to recover from catastrophic events. However, our increasingly non-stationary world is also generating an unprecedented amount of raw data, and the tools and techniques developed by the data science community are significantly reducing the costs of extracting signals and patterns that allow for a more dynamic view of risk. This presentation discusses two perils, wildfires and cyber risks, where new data streams can provide insight into trends towards increasing catastrophic losses. Nine out of the ten most damaging wildfires in U.S. history have occurred since 2017, driven by dry vegetation, high winds, and increased development in the urban-wildland interface. Modern remote sensing tools generate large amounts of geospatial data that can be used to understand how large and fast fires might grow under a range of weather, vegetative, and topographic conditions. Likewise, recent catastrophic cyber events demonstrate how modern ransomware is increasingly able to exploit interconnected software systems to cause systemic (and costly) disruption. The cybersecurity community generates a large quantity of textual data about breach events and potential vulnerabilities meant to inform mitigation and alert users to risk. In both cases, insurers lack the ability to translate these data streams into probabilistic, interpretable metrics that can be used by actuaries and risk managers. Efforts to develop new methods of quantifying risk and opportunities for deeper exploration are discussed.

Self-supervised In-context Operator Learning on the Probability Measure Space

Rongjie Lai

Purdue University, USA

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 a self-supervised 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.

O(k)-Equivariant Dimensionality Reduction on Stiefel Manifolds

Harlin Lee

University of North Carolina at Chapel Hill, USA

Many problems in science involve data that live on matrix manifolds, where respecting geometric structure and symmetry is essential for meaningful representation learning. In this talk, we address the problem of dimensionality reduction for data on Stiefel manifolds while preserving invariance under the orthogonal group O(k). We illustrate the approach with numerical experiments on neuroscience and video data and discuss extensions to Grassmannian manifolds.

Seeing the Invisible with Medical Imaging and Data Science: Challenges and Opportunities

Yueh Lee

University of North Carolina at Chapel Hill, USA

Radiology is a common focus of conversations about the future of AI, especially when applied to medicine. Medical imaging is an integral part of medical care, and thus how AI could be applied in medicine. Currently, 75% of the current FDA-approved AI devices are for the support of the field of Radiology, but adoption remains slow. The volume of X-rays, CTs, MRIs, Ultrasounds, and molecular imaging offers a potentially rich data set for data science, but presents potential challenges for the uninitiated. The goal of this talk is to explain the challenges associated with medical images, review current data science / AI approaches, and offer ideas about future paths.

AI for Scientific Discovery: Learning Beyond Data in Computational Wave Imaging

Youzuo Lin

University of North Carolina at Chapel Hill, USA

Artificial intelligence is transforming how we explore and understand the natural world. In scientific domains such as computational wave imaging, from mapping Earth's subsurface in geoscience to visualizing tissues in medical ultrasound, AI provides powerful new tools to model, reconstruct, and interpret complex physical phenomena. Yet, a fundamental question remains: how can we design learning systems that capture the governing physics of a process, rather than merely memorizing data? In this talk, I will discuss two complementary directions toward this goal. The first emphasizes physics-informed learning, where physical laws are embedded directly into model architectures or training objectives, enabling robust inference even with limited data. The second examines the rise of large-scale and foundation models, which leverage massive simulated or experimental datasets to learn broad, transferrable representations across physical regimes. By contrasting these paradigms, physics-guided versus data-scaled, I will highlight their respective strengths, limitations, and potential for hybrid integration. Drawing on examples from my research in geophysical and biomedical wave imaging, I will illustrate how these approaches can enhance both understanding and practical performance. I will conclude with a broader perspective on AI-driven scientific discovery, focusing on how teaching machines the principles of physics can move us beyond data fitting toward genuine understanding of the physical world.

A General Framework for Group Sparsity in Hyperspectral Unmixing Using Endmember Bundles

Yifei Lou

University of North Carolina at Chapel Hill, USA

Due to low spatial resolution, hyperspectral data often consists of mixtures of contributions from multiple materials. This limitation motivates the task of hyperspectral unmixing (HU), a fundamental problem in hyperspectral imaging. HU aims to identify the spectral signatures (endmembers) of the materials present in an observed scene, along with their relative proportions (fractional abundance) in each pixel. A major challenge lies in the class variability in materials, which hinders accurate representation by a single spectral signature, as assumed in the conventional linear mixing model. To address this issue, we propose using group sparsity after representing each material with a set of spectral signatures, known as endmember bundles, where each group corresponds to a specific material. In particular, we develop a bundle-based framework that can enforce either inter-group sparsity or sparsity within and across groups (SWAG) on the abundance coefficients. Furthermore, our framework offers the flexibility to incorporate a variety of sparsity-promoting penalties, among which the transformed L1 (TL1) penalty is a novel regularization in the HU literature. Extensive experiments conducted on both synthetic and real hyperspectral data demonstrate the effectiveness and superiority of the proposed approaches.

Expectation-enforcing Strategies for Repeated Games

Alex McAvoy

University of North Carolina at Chapel Hill, USA

The theory of repeated games describes interactions that unfold over time. Strategies for repeated games can condition on the past in complex ways, and it is known that this behavioral complexity yields rich spaces of equilibria. However, it often results in an equilibrium selection problem, and in many interesting games there isn't a reliable way to get to "good" outcomes. In this talk, I will describe a class of "autocratic" strategies that allow a player to unilaterally control the space of possible outcomes. Examples include fairness (e.g., my payoff equals yours, on average) and extortion (e.g., my payoff equals twice yours, on average). These strategies have nice adaptive properties in the sense that they incentivize a self-interested opponent to also act in the interests of the group. Within this class of strategies, I will describe how an agent can always find an autocratic strategy with short memory, which takes into account information from only one round into the past. In particular, we can characterize precisely which payoff relationships are enforceable using strategies of arbitrary complexity. This memory reduction indicates that certain kinds of reciprocal payoff control might be feasible for humans (who are notoriously bad at synthesizing information over long time horizons). It also demonstrates the sufficiency of a short-memory search space for autocratic strategies in reinforcement learning.

Steering Large Language Models: A Geometric and Control-Theoretic Approach

Minh Tan Nguyen

National University of Singapore, Singapore

Inference-time activation steering offers a lightweight way to control foundation-model behavior without updating weights, but standard "add-a-direction" interventions can be brittle and hard to tune. This talk highlights two first-principles advances from my group for more reliable control: Angular Steering, which treats steering as a geometric rotation in a low-dimensional subspace to enable smooth, continuous control with a single interpretable knob; and PID Steering, which frames behavior control through classical feedback control, adding integral and derivative terms to improve robustness, reduce drift, and stabilize outcomes across prompts and settings. Together, these methods make behavioral control more predictable, tunable, and principled for modern foundation models.

Gromov-Wasserstein Problem: Non-convexity, Semidefinite Relaxations, and Gradient Flows

Yongsheng Soh

National University of Singapore, Singapore

The Gromov-Wasserstein (GW) problem is an extension of the (Kantorovich formulation of the) Optimal Transport problem to settings where a cost function that describes the price of moving mass from one location to another is absent. Instead of finding optimal couplings that minimize expected transportation cost, the GW seeks couplings that minimize the distortion between metric spaces associated with the source and target space. Unfortunately, the formulation, as it is stated, has quadratic dependency on the coupling variable and is non-convex, which introduces computational challenges involving global optimality that is not encountered in classical Optimal Transport. In this talk, I will describe a semidefinite relaxation for the GW problem. The derivation of the relaxation is based on similarly developed relaxations for solving the closely associated graph matching problem and the Quadratic Assignment Problem. The semidefinite relaxations always outputs a feasible coupling plan. In particular, our method provides a principled manner of certifying if a computed plan is globally optimal. In fact, our numerical experiments on moderate-sized problems with real data show that our relaxations often succeed at finding the globally optimal solution, which suggests that the proposed relaxation is indeed strong. Finally, we consider the problem of computing gradient flows with respect to the GW metric. We propose a simple computational procedure for computing such flows as metric spaces that we later embed in Euclidean space, which allows one to obtain explicit visualizations of the evolution of metric measure spaces with respect to the GW metric.

Adapting Noise to Data by Quantile Learning

Gabriele Steidl

Technische Universität Berlin, Germany

The default Gaussian latent in flow-based generative models poses challenges when learning certain distributions such as heavy-tailed ones. We introduce a general framework for learning data-adaptive latent distributions using one-dimensional quantile functions, optimized via the Wasserstein distance between noise and data. The quantile-based parameterization naturally adapts to both heavy-tailed and compactly supported distributions and shortens transport paths. Numerical results confirm the method's flexibility and effectiveness achieved with negligible computational overhead.

Wasserstein Gradient Flow of Maximum Mean Discrepancy with Energy Kernel

Lihan Wang

National University of Singapore, Singapore

Wasserstein gradient flows of Maximum Mean Discrepancy have recently found applications in statistics and machine learning. The standard theory of Wasserstein gradient flows in spaces of probability measures provides well-posedness for smooth kernels, however, the kernels that perform well in applications are not differentiable and the theory does not apply. Here we consider the "energy kernels" $k(x, y) = -|x - y|^q$ for $q \in (0, 2)$ which allow for long-range effects. We develop well-posedness theory in L^p spaces and show that these solutions can be approximated well by interacting particles, which justifies their use. We also provide results that show that at least in 1D the convergence can be exponential when the solution is close to the target and also prove that no quantitative rate of convergence holds for general initial data. Joint work with Matthew Rosenzweig and Dejan Slepčev (CMU).

Towards Large Scientific Learning Models with In-Context Operator Networks (ICON)

Liu Yang

National University of Singapore, Singapore

Can we build a single large model for a wide range of scientific problems?We proposed a new framework for scientific machine learning, namely “In-Context Operator Learning” and the corresponding model “In-Context Operator Networks” (ICON). A distinguishing feature of ICON is its ability to learn operators from numerical prompts during the inference phase, without weight adjustments. A single ICON model can tackle a wide range of tasks involving different operators, since it is trained as a generalist operator learner, rather than being tuned to approximate a specific operator. This is similar to how a single Large Language Model can solve a variety of natural language processing tasks specified by the language prompt. We will show how a single ICON model (without fine-tuning) manages multiple distinct problem types, encompassing forward and inverse ODE, PDE, and mean-field control problems. Through a case study on 1D conservation laws, we will show ICON’s strong generalization capability to new PDEs, as well as its advantage compared with classic operator learning methods, e.g., Fourier neural operator (FNO). We will also show the application of ICON in 2D fluid problems, where a single model can make predictions for incompressible or compressible fluids, with different viscosity.

Accelerate Neural Operators by Quantum Computing

Xiu Yang
Lehigh University, USA

In the realm of computational science and engineering, constructing models that reflect real-world phenomena requires solving differential equations with different conditions. Recent advancements in neural operators, such as deep operator network (DeepONet), which learn mappings between infinite-dimensional function spaces, promise efficient computation of differential equation solutions for a new condition in a single forward pass. However, classical DeepONet entails quadratic complexity concerning input dimensions during evaluation. Given the progress in quantum algorithms and hardware, here we propose to utilize quantum computing to accelerate DeepONet evaluations, yielding complexity that is linear in input dimensions. We demonstrate the method's efficacy in both ideal and noisy conditions. Furthermore, we show that our quantum DeepONet can also be informed by physics, minimizing its reliance on extensive data collection. Quantum DeepONet will be particularly advantageous in applications in outer loop problems which require exploring parameter space and solving the corresponding differential equations, such as uncertainty quantification and optimal experimental design.

What can One Expect when Training Neural Networks to Solve PDEs?

Hongkai Zhao

Duke University, USA

When neural networks (NNs) are used as a type of nonlinear parametric representation to solve partial differential equations (PDEs), they often display frequency-dependent learning dynamics that can differ from those seen in direct function approximation tasks, resulting from a balance between the frequency bias of the NN representation and that of the underlying differential operator. Using elliptic PDEs as examples, we show how these two factors compete and balance in different situations. Once the balance is determined, it is important to design computational strategies to counter the resulting bias to improve training efficiency. We propose a simple operator aware preconditioning strategy that rebalances the optimization landscape and the learning dynamics. Extensive experiments, including multiscale problems, show that our approach restores more balanced learning dynamics across modes and substantially improves both convergency and accuracy.