Caltech
Student-Faculty Programs

Session G Abstracts

Neural operators for dissipative relativistic magnetohydrodynamics
Ansh V. Desai
Mentors: Anima Anandkumar, Chuwei Wang, Elias R. Most, and Valentin Duruisseaux

Relativistic dissipative magnetohydrodynamics is essential for modeling high-energy astrophysical systems such as black hole accretion flows, where non-equilibrium effects like shear stress, heat flux, and magnetic conduction play a critical role. While these dissipative contributions are necessary for accurate simulations, their repeated evaluation in large-scale numerical solvers is prohibitively expensive, often requiring hundreds of millions of GPU hours. In this work, we utilize a Fourier Neural Operator based framework to bypass this computational bottleneck by learning a surrogate model for the evolution of dissipative quantities from easily computed local fluid states. This approach takes advantage of the remarkable discretization-invariance and generalization capabilities of neural operator to predict small-scale dissipative dynamics with high efficiency, enabling rapid integration into global simulations.

Shape optimization with neural operators
Jiayi Zhou
Mentors: Anima Anandkumar and Valentin Duruisseaux

Shape optimization is central to modern engineering design. A typical pipeline involves shape parameterization, PDE-based evaluation of performance metrics, and iterative optimization. In this work, we propose a fully differentiable shape optimization framework that leverages neural networks (NNs) for shape representation and neural operators (NOs) for efficient PDE approximation. Unlike traditional approaches such as control nodes, which require extensive manual setup and offer limited flexibility, our NN-based parameterization supports complex geometries and enables gradient-based optimization. We also aim to address the incorporation of geometric constraints. Another key focus of our work is numerical robustness and convergence behavior during optimization. We investigate the distribution-invariance property of neural operators, ensuring consistent performance across varying discretizations and point cloud samplings.

LeanLibrary: A unified framework for theorem proving in Lean 4
Ryan L. Hsiang
Mentors: Anima Anandkumar and Robert Joseph George

Mathematics is experiencing a fundamental transformation through the integration of artificial intelligence and formal verification. Despite the significant advancements in combining large language models with the Lean theorem prover, much work is needed to address the critical challenges in codebase maintainability and performance limitations affecting user experience. This project integrates the various Lean+AI tools into a unified library, providing a consistent API for data extraction, interaction, model fine-tuning, and AI-assisted theorem proving. In addition, our work involves reimplementing the training pipeline to support modern transformer architectures, extending LeanCopilot to allow external API-based inference with GPU acceleration, and building Lean4Code, a user-friendly IDE for mathematicians.

OrbNet-Materials: Orbital-based GNN for materials properties prediction
Minhyuk Kang
Mentors: Anima Anandkumar and Beom Seok Kang

Accurate electronic band structures are fundamental to materials discovery, but DFT is computationally expensive and most ML surrogates either rely on geometry-only inputs or predict only scalar properties rather than the full band energy. Our OrbNet-Materials, a periodic extension of orbital-learning GNN can predict full electronic band structures for crystalline materials at a fraction of DFT cost. Building on the success of OrbNet in orbital learning for molecular system, we generate

k-dependent atomic-orbital matrices via DFTB+ using low-cost semiempirical quantum mechanical (QM) calculations (GFN1-xTB) and feed them to an SE(3)-equivariant OrbNet framework. A PBC-aware graph links each unit cell to intra-cell and periodic-image neighbors within a cutoff, and per-*k* inference predicts $\varepsilon_n(k)$ and derived properties. We train with a dispersion-aware loss: pointwise energy MSE plus finite-difference penalties on band slops, with optional extra weight at high-symmetry points. Together, these components establish an end-to-end pipeline for orbital learning in periodic materials. OrbNet-Materials has strong potential to serve as a fast, orbital-aware surrogate for DFT in high-throughput discovery, driving inverse design and active-learning loops for photovoltaics, thermoelectrics, and catalysis.

Partially observable model-based reinforcement learning for drag reduction in compressible turbulent flows
Taeyang Park
Mentors: Anima Anandkumar and Myrl Marmarelis

Active control of compressible turbulent flows presents significant challenges not found in incompressible regimes, including long-term instabilities such as acoustic waves and local shocklets that destabilize simple control strategies. This work presents a framework for Partially Observable Model-Based Reinforcement Learning to address this challenge, evolving the PINO-PC architecture by leveraging a differentiable CFD solver (JAX-Fluids). We augment the original observer model to function as a multi-step "world model" that is trained on the full state information from the simulation. This augmented observer is designed to perform stable "rollouts", enabling the control policy, which operates on realistic partial observations from wall-mounted sensors, to anticipate and counteract the slow-building instabilities that myopic, single-step predictions would miss. We detail the architectural modifications required for this approach and ensuring the stability of the learned world model in multi-step predictions.

Discovery and parameter estimation of PDEs using physics-informed neural networks (PINNs)
Edgar A. Larios
Mentors: Franca Hoffmann, Aras Bacho, and Kathrin H. Hellmuth

Jupiter captures high radiation levels and large quantities of radiation, making it very difficult for NASA spacecraft to probe the planet or its surroundings. To better design its missions, NASA relies on models that describe the radiation environment of Jupiter, as well as data collected from previous missions. The radiation environment of Jupiter can be represented by a partial differential equation (PDE) that is structurally similar to the Fokker-Planck equation. A physics-informed neural network (PINN) was used to discover, solve, and invert this PDE. Some of the techniques that were used to create a robust PINN-based pipeline include data normalization and hyperparameter tuning. To verify the reliability of the PINN, it was applied to synthetic benchmark problems. Results show that the PINN was able to solve all the PDEs (including the PDE that describes the radiation environment of Jupiter), but was always unable to properly recover the coefficients of the PDEs. These findings suggest that PINNs can solve forward problems reliably, but it is difficult to make them suitable for inverse problems. I recommend continued testing on synthetic benchmark problems to improve the robustness of the developed PINN-based pipeline before applying it to study Jupiter’s radiation environment.

Optimal experimental design for Jupiter's radiation belt
Cristian D. Peña
Mentors: Franca Hoffmann, Aras Bacho, and Kathrin H. Hellmuth

Understanding Jupiter’s radiation environment is critical for designing safe and effective spacecraft missions, yet the limited availability of observational data makes their dynamics mostly unknown. This project investigates how mathematical tools can improve the scientific return of spacecraft missions aimed at studying Jupiter’s radiation belt. We first model Jupiter’s radiation belts using a partial differential equation, the Fokker–Planck equation, whose coefficients must be estimated from sparse data. To estimate these coefficients, we pose a Bayesian inverse problem that incorporates prior knowledge and quantifies uncertainty while inferring model inputs from data. By parameterizing

candidate paths and maximizing the determinant of the Fisher Information Matrix, which quantifies trajectory informativeness, we then identify spacecraft paths that offer the greatest potential for reducing model uncertainty. These efforts are supported by the adjoint method, which enables efficient computation of how changes in the model's coefficients influence both its predictions and the informativeness of different designs.

A survey of data-driven techniques for network inference
Owen M. Tolbert
Mentors: Andrew M. Stuart and George Stepaniants

Networks are a crucial component of research across various scientific disciplines, and inferring the connectivity of those networks with sparse observational data is a pertinent challenge. Here we explore methods from causal inference and machine learning to solve this problem. We begin by studying the Synergestic, Unique, Redundant, Decomposition (SURD) framework to causally infer the connections between nodes in a networked dynamical system. In this direction, our results are inconclusive as the SURD framework does not infer unique dependencies between nodes and does not scale to larger network sizes. Next we explore the use of neural network architectures, specifically variational autoencoders, to recover network connectivity from time series data. Our preliminary experiments apply these methods to networks of mass-spring and Kuramoto oscillators, showing how they infer the structure of linear as well as nonlinear dynamical systems.

Blow-up scenarios in the Keller-Segel system
Zirui Wang
Mentors: Thomas Y. Hou and Xiang Qin

Blow-up phenomena are central to the study of nonlinear partial differential equations, capturing finite-time singularity formation and revealing deep links between analytic structure and dynamics. A key question is whether a solution admits a type I or type II blow-up ansatz, as this distinction governs both the qualitative behavior and the analytical approach. In many cases, the absence of an explicit blow-up profile makes determining the rate substantially more challenging.

In this work, we focus on the Keller-Segel system, including both the variant with logistic damping and the high-dimensional case. We investigate the structure of exact self-similar profiles and the possibility of type II blow-up through a combination of asymptotic analysis and high-resolution numerical simulations, providing quantitative blow-up rate estimates and insights into the stability of the underlying profiles.

Kernelized Stable Fluids for simulating physically accurate solutions
Lennart A. Scholz
Mentors: Houman Owhadi and Aras Bacho

Solving the Navier–Stokes equations is of great importance both in physics and applied mathematics. Due to their nonlinearity, most of the established solution algorithms, such as the pseudospectral method, have shortcomings, including stability problems for large timesteps. This problem is addressed by the Stable Fluids method, which solves the Navier–Stokes equations in an unconditionally stable manner. However, the method has a disadvantage: the solution of the advection step is not volume-preserving. The method solves the advection part by tracing each point backwards in time along the divergence-free velocity vector field of the fluid. The flow of a divergence-free velocity vector field is theoretically volume-preserving, but not if we approximate it using the method of characteristics, as proposed by the Stable Fluids method. We propose Kernelized Stable Fluids in order to ensure volume preservation when solving the advection step, thereby making the algorithm suitable for scientific simulations and not just visual effects.

Coarse-to-fine diffusion language models

Frank Y. Xiao
Mentors: Pietro Perona and Rogerio Aristida Guimaraes

Diffusion Language Models (DLMS) have recently emerged as an exciting potential alternative to autoregressive models. However, the current masked diffusion paradigm does not take advantage of the intermediate computations done during the diffusion reverse process, instead collapsing every masked token to a single generic mask during every intermediate step. We propose a coarse-to-fine diffusion process that introduces semantically meaningful intermediate tokens that allow DLMs to pass tokens through a hierarchy of masks before settling on their final lexical values. We (i) define a new discrete forward noising process that respects arbitrary mask hierarchies, (ii) generalize the Rao-Blackwellised ELBO to multi-transition chains, yielding a low-variance training objective, and (iii) design an efficient ancestral sampler that exploits the hierarchy for faster decoding. We benchmark our methods on generative perplexity tasks on LM1B and explore the viability of coarse-to-fine DLMs.