

Multi-objective optimization for synthetic aviation turbine fuels

Marissa Till

Mentors: Steven Brunton, John O. Dabiri, Ana Larrañaga, and Patricia Garcia Caspueñas

Synthetic Aviation Turbine Fuels (SATFs) offer environmentally friendly alternatives to conventional kerosene-based jet fuels, but identifying optimal blends is a largely unexplored multi-objective design problem. Fuel blends must balance competing goals—such as reducing contrails and minimizing pollutant emissions—while satisfying nonlinear chemical blending rules, Federal Aviation Administration regulations, and engine compatibility requirements. In this work, SATF development is framed as an inverse design problem to identify chemical compositions that meet performance targets defined by a user-specified merit function. Building on a baseline single-objective, weighted-sum Genetic Algorithm (GA) that minimized Landing-and-Takeoff (LTO) emissions, a refined framework was introduced that couples global, evolutionary multi-objective search algorithms, such as Non-dominated Sorting Genetic Algorithm III (NSGA-III) and Multi-Objective Evolutionary Algorithm based on Decomposition (MOEA/D), with local refinement methods, including Nelder–Mead simplex and Constrained Optimization by Linear Approximations (COBYLA). While exploring trade-offs between soot and pollutant emissions, these methods conduct constrained optimization of fuel blend properties across blends of ten to twelve chemical species selected from a database spanning five chemical families. These multi-stage pipelines demonstrate improved optimization robustness while generating a diverse Pareto front of optimal fuel blends, enabling manufacturers to interpret trade-offs and select fuels aligned with their operational priorities and regulatory constraints.

Sparse identification of nonlinear dynamics via Large Language Models (SINDy-LLM)

Atharv Prabhutendolkar

Mentors: Steven Brunton and Nick R. Hutzler

Discovering the governing differential equations of physical systems from noisy, observational time-series data is a fundamental challenge in scientific modeling. While traditional system identification methods—such as the Sparse Identification of Nonlinear Dynamics (SINDy)—offer computational precision, they are constrained by rigid candidate libraries, suffer from noise amplification, and lack domain-specific physical reasoning. Conversely, Large Language Models (LLMs) possess vast scientific prior knowledge but are severely limited by computational inaccuracies and mathematical hallucinations.

To bridge this gap, we present SINDy-LLM, an autonomous, closed-loop agentic framework that couples the qualitative reasoning of LLMs with a suite of specialized numerical solvers. SINDy-LLM introduces a dynamic Retrieval-Augmented Generation (RAG) diagnostic troubleshooting loop: when numerical fitting yields poor fits ($R^2 < 0.95$), the agent automatically triggers diagnostic playbooks to pre-process data or switch solvers. Additionally, we integrate a decoupled parameter optimization tool (FitSkeleton) that enables the discovery of non-polynomial and transcendental equations (e.g., trigonometric forces in pendulums) by fitting LLM-proposed symbolic structures using gradient-based optimizers.

We evaluate SINDy-LLM across a diverse benchmark of non-linear and chaotic systems (including the Duffing oscillator, Lotka-Volterra, and Van der Pol limit cycles) under 5% measurement noise. The agentic tool-calling loop achieves up to a 2,800x reduction in trajectory simulation error (NMSE) compared to standard zero-shot LLM predictions. Finally, we establish an end-to-end instruction fine-tuning pipeline, scaling a custom dataset to 28 multi-turn scientific dialogues to train a domain-specific "SINDy Scientist" model. Our results demonstrate that combining LLM reasoning with automated numerical tool-use represents a robust, highly generalizable paradigm for autonomous scientific discovery.

Python port and evaluation of immersed boundary projection method

Yixuan Lin

Mentors: Steven Brunton, Soon-Jo Chung, Ana Larrañaga, and Patricia Garcia Caspueñas

Prediction and control of fluid flows are one of the most challenging problems in engineering. The Immersed Boundary Project Method (IBPM) is a promising numerical approach for resolving incompressible flow evolution around complex, changing geometries. IBPM uses a multi-domain structured grid to initially simulate the flow field and subsequently correct for fluid-solid interactions, eliminating the need for time-consuming mesh-generation and accelerating computation. This work ports the IBPM to Python, positioning it as a potential candidate for Hydrogym 2.0, a fluid environment platform for the application of reinforcement learning in flow control. The code implementation is complemented by a thorough physical analysis of the method's limitations. We study the solver across different airfoil geometries and compare the resulting velocity fields against established data from alternative numerical methods.

AutoMF: Automated ensembles for multifidelity field prediction

Eloise Zeng

Mentors: Faez Ahmed, Tony Yue Yu, and Nicholas Sung

Predicting fine-resolution physical fields from simulation parameters is costly, and no single surrogate performs best across tasks. Multifidelity learning combines coarse and fine solutions, but selecting a single surrogate from limited fine data can discard useful predictions. We present AutoMF, a workflow that uses reserved fine solutions to select or combine trained field predictors. Its library contains nine neural and reduced-basis models that use coarse and fine training fields in different ways, including transfer learning, pooled-fidelity training, coarse-field correction, and a fine-only reference. For each dataset, five additional parameter-field pairs determine either a selected model or fixed ensemble weights based on individual or combined prediction errors. These weights remain fixed across new inputs and grid points, and inference requires no additional coarse simulation. We evaluate AutoMF on 22 datasets spanning 21 partial differential equation tasks in seven classes and one climate-emulation task. Using the same five fitting examples, the optimized mixture improves on the model selected from those examples on 17 of 22 datasets. Relative to this selected model, it reduces geometric-mean error by 11.6% when PDE classes receive equal weight and by 7.5% when all datasets receive equal weight. Inverse-error weighting achieves most of the aggregate gain. These results show that combining heterogeneous field predictors can reduce the cost of model-selection errors when the best surrogate for a new task is unknown.

Manipulating the domain of Rayleigh-Taylor turbulence: Confined vs. self-similar RT regimes

Riley E. Foard

Mentor: Guillaume Blanquart

Dense fluid on top. Thin fluid on bottom. Add gravity, wait a moment, and chaos ensues. Rayleigh-Taylor (RT) instabilities begin with a simple density difference but quickly develop into a turbulent mixing layer. Two limiting regimes govern a temporally-growing RT (TRT) mixing layer. In the self-similar regime, the layer grows according to the familiar scaling $h \propto t^2$. However, once the mixing layer grows much thicker than the horizontal domain, the flow then enters the confined regime, scaling as $h \propto t^{2/5}$. This research investigates that transition while asking a second question: can both regimes be recreated without continuously growing the simulation? Statistically stationary Rayleigh-Taylor (SRT) simulations offer a computational shortcut, holding mixing layer height fixed while turbulence continues to evolve. SRT incorporates growth assumptions into forcing terms of the transformed Navier-Stokes equations, prescribed independently of geometry. This research derives these terms and isolates the effects of growth scaling and geometry. Using direct numerical simulations, TRT and SRT flows are compared across various metrics, including mean mole fraction, kinetic energy, anisotropy, and local growth exponent. Ultimately, this research accesses the same long-time physics without a larger domain, longer simulations, and an increasingly painful computational bill.

Impact of turbulent inflow treatment on detonation-turbulence interaction

Max R. Ismagilov

Mentors: Guillaume Blanquart and Benjamin R. Breer

Direct numerical simulation (DNS) of shock- and detonation-turbulence interaction (DTI) relies on several methods of generating turbulent inflow. The first and most widely used method is Taylor's hypothesis, whereby a frozen spatial realization of turbulence is advected through the detonation at the Chapman-Jouguet (CJ) propagation speed. This method can introduce artificial periodicity, as the length of the inflow is shorter than the residence time in the detonation. A second and more computationally expensive option is to perform a separate DNS of compressible homogeneous isotropic turbulence (HIT) with a mean convective speed in the streamwise direction (equal to the CJ speed); transverse planes of velocity are recorded as a function of time, resulting in a spatially and temporally dynamic turbulent inflow signal. Typically, the triply periodic HIT domain is cubic in shape; hence large scales (convecting at the shock propagation speed) are "recycled" as they repeatedly pass through the sampling plane before becoming decorrelated over one integral time scale. Thus, a third method for generating turbulent inflow uses an elongated HIT domain, wherein the residence time is greater than the large eddy turnover time. Finally, a fourth method simply uses solenoidal velocity fluctuations (from incompressible HIT) as the turbulent inflow; however, it remains unclear whether the injected solenoidal velocities would yield the appropriate dilatational fields before turbulence reaches the leading shock. In this study, each of these four methods of inflow generation are performed for a single DTI case with fixed upstream turbulence conditions. Statistics (means, planar averages, probability density functions of shock speed and induction length) are computed and compared across the different methods of inflow generation.

Physics-based versus learned state-space models for building envelope thermal dynamics under passive excitation

Brendan J. Rudberg

Mentors: Ólafur P. Pálsson and Dependra Gautam

Stochastic state-space models span a spectrum of inductive bias, from gray-box resistor-capacitor (RC) networks, whose structure is governed by physics and whose parameters are estimated by maximum likelihood via Kalman filtering, to learned architectures such as long short-term memory networks (LSTMs), which infer dynamics directly from data. We use a single instrumented building envelope, with air- and wall-surface-temperature sensors, to compare models across this spectrum on a common basis. We identify a gray-box wall model by forward selection over increasingly complex RC networks, using air temperatures and solar radiation as inputs and wall temperatures as observations. AIC selects a two-capacitance model with a solar-gain term; higher-order networks fail to converge. We benchmark this model against a graded lineup of persistence and static-linear baselines, a linear black-box state-space model, and an LSTM. This progression separates the contributions of exogenous inputs, temporal dynamics, and learned nonlinear flexibility, allowing held-out multi-step prediction error on a chronological split to be compared across modeling assumptions. The central question is whether an explicit physics prior improves prediction when data are limited and passively excited. With 33 days of data and no controlled input, the seven-parameter gray-box model matches the 4,802-parameter LSTM on the outer surface (0.98 vs. 0.99 °C RMSE), while static regression and the LSTM outperform it on the near-static inner surface. These results show that an explicit physics prior can provide competitive generalization under limited, passively excited data, particularly for the more dynamically informative outer surface, though its parameters could not be reliably estimated from the passive data.

Designing rigid-soft hinges for microscale Kresling origami metamaterials for future application in micro robotics

Evan W. Dickson

Mentors: Glaucio Paulino and Chiara Daraio

This work investigates strategies for achieving rigid-soft mechanical contrast in microscale Kresling origami metamaterials. Kresling structures rely on coupled axial-rotational deformation and bistability, but microscale implementations often fail through hinge fracture, bar buckling, or unintended

deformation because fabrication cannot reproduce the rigid-panel/compliant-hinge assumption of ideal origami models. Reduced-order Kresling mechanics are used to identify the bar-to-hinge stiffness ratio and allowable hinge strain as primary design constraints. Three fabrication strategies are evaluated: exposure-programmed two-photon polymerization, step-growth photopolymerization with a tunable modulus of approximately 7 MPa to 1.3 GPa, and selective metal coating of load-bearing regions. The analysis indicates that material contrast alone may be insufficient and should be combined with flexure geometry and selective bar stiffening to preserve bistability without exceeding hinge yield limits. This framework establishes a pathway for designing, fabricating, and characterizing reliable microscale Kresling structures for programmable metamaterials and remotely actuated microrobotic systems.

Dexterous five-fingered tendon-driven anthropomorphic robotic hand platform implementing a field-oriented control actuation architecture

Jilly A. Choi

Mentor: Aaron D. Ames

Dexterous robotic hands for humanoids require both high force output and mechanical robustness. Conventional robotic hand architectures often rely on high-reduction gearboxes to amplify the force of embedded motors. However, this motor actuation architecture limits backdrivability and impact tolerance due to the high reflected inertia of the gear transmission. This work presents the development of a five-fingered, tendon-based anthropomorphic robotic hand platform incorporating a field-oriented control actuation architecture combined with a backlash-free force-amplifying pulley mechanism. These mechanisms increase mechanical leverage to drastically increase fingertip force without adding backlash or reflected inertia. The FOC architecture utilizes motor current for direct torque regulation, enabling highly backdrivable, mechanically transparent actuation that redirects external impact loads into backdriving a motor, directly converting kinetic energy into electrical energy that is dissipated as heat. Additionally, the mitigation of backlash allows for real-time proprioception of the fingers' external forces through in-line current measurement from the brushless motors. The finger architecture implements three motors to control the distal and proximal interphalangeal joint "curl," metacarpophalangeal joint "rotation," and abduction/adduction through a differential pulley mechanism designed within the MCP joint. Furthermore, the design establishes a platform for future comparison and performance characterization against conventional robotic hand architectures.

Improving object permanence under occlusion in latent video prediction for robot manipulation

Yiling Hou

Mentor: Aaron D. Ames

World models that capture physical dynamics, rather than only appearance, are increasingly seen as a foundation for robot control. A basic test of physical understanding is object permanence: an object hidden from view still exists and keeps moving. We study this in V-JEPA 2, a Joint-Embedding Predictive Architecture (JEPA) that predicts future video in latent space. In robot manipulation, targets are often briefly occluded by the robot's own arm or by a held object. Using controlled simulations, we find that the model's latent representations of hidden objects degrade during occlusion, especially for objects in motion. A targeted adaptation of the predictor substantially reduces this degradation. We are now evaluating whether this physics-aligned predictor improves a vision-language-action (VLA) policy on a robotic task in which a moving target is temporarily hidden from the camera.

Integrated CFD analysis and control design for jet interaction in interceptors

Rajat V. Bidarkota

Mentors: Karen M. Siegrist, Joanna M. Austin, John O'Neil, Nicholas Fezie, and Walter Maier

Hypersonic glide vehicles can maneuver during descent in ways that traditional ballistic missile defense systems struggle to counter. Meeting the resulting high-altitude maneuvering requirements motivates interceptors that steer using lateral gas jets instead of aerodynamic control surfaces, a mechanism known as jet interaction. However, jet interaction introduces complex, geometry-dependent flow physics that must be characterized before it can be used reliably for guidance and control. This project uses computational fluid dynamics to characterize jet interaction on a generic

interceptor over a range of thrust levels, angles of attack, and roll angles, from which force and moment amplification factors are extracted. These factors are incorporated into a six-degree-of-freedom guidance and control simulation by modifying an existing divert-and-attitude-control system to account for jet interaction effects. Preliminary results indicate significant force amplification at low jet thrust momentum ratios relative to freestream dynamic pressure, while aerodynamic angles and fin geometry can substantially amplify or diminish net thrust depending on nozzle orientation and shock interaction, resulting in a modeled vehicle that is difficult to stabilize and maneuver. Ongoing work focuses on achieving stable control against maneuvering hypersonic targets and optimizing thruster placement and control logic to minimize miss distance.

Visual anemometry in urban environments using Ned Kahn's *Shimmer Wall* kinetic sculpture

Delta M. Blendea

Mentors: Nathaniel Wei and Gil Refael

Winds in urban environments dictate the dispersal of pollutants, the formation of heat islands, and provide natural ventilation. Despite their importance for human flourishing, relatively little is understood about the nature of urban winds at city-block scale. The Shimmer Wall, a kinetic facade by artist Ned Kahn, visualizes the motions of urban winds using swinging aluminum panels, thus acting as a large-scale visual anemometer. By observing the sculpture with video cameras, we can determine the magnitude, spatial scale, and advection velocity of pressure fluctuations. Lab experiments performed both indoors and outdoors with a scale model of the wall demonstrate that regardless of lighting conditions, a change in panel brightness is correlated with a change in panel angle. As the panels largely obey the differential equation for a forced, damped pendulum, we derive a transfer function to map panel angular velocity to forcing torque, providing the normal stress and pressure differential across the panel. This analytical solution complements a numerical method to solve for the same using a normalized least-squares optimizer. Across over 10,000 panels spanning 20 meters horizontally and 15 meters vertically, the data provides a time-resolved, large-scale measurement of urban wind effects. These results demonstrate the feasibility of low-cost, non-intrusive volumetric wind measurements in urban environments.