

2026 Student Abstract Book



Frequency of Small Magellanic Cloud analogs in simulated Milky Way systems hosting a Large Magellanic Cloud analog

Selassie A. Ackah

Mentors: Andrew J. Benson and Lior S. Pachter

The Milky Way anchors studies of galaxy formation and dark matter, yet its satellite system may be unusual because it contains both Magellanic Clouds. We investigate how often a Small Magellanic Cloud analog accompanies a Large Magellanic Cloud analog in simulated Milky Way systems. From 10,000 Galacticus merger-tree realizations, we selected candidates with subhalo masses of 5-8 billion solar masses, three-dimensional separations of 20-40 kiloparsecs, and evidence of prior association with the Large Cloud analog. In the current analysis, this definition identifies 618 candidates in 606 systems, a preliminary candidate-hosting fraction of 6.06%. Removing the association requirement raises the fraction to 27.65%, indicating that mass and distance alone select many objects lacking a shared accretion history. These estimates remain subject to validation of the mass definition, association criterion, and selection pipeline. Nevertheless, inferred Small Magellanic Cloud frequency is highly sensitive to how physical association is defined. Robust criteria are necessary for judging whether the Milky Way is representative and for testing galaxy-formation and dark-matter models.

Inferring supernovae progenitor properties from a systematic sample of nebular spectra

Athiwat N. Ackaramongkolrotn

Mentors: Mansi M. Kasliwal and Kaustav Das

Determining the progenitor masses of supernovae are crucial to understanding the processes that underlay their formation. Of particular interest are stripped-envelope core-collapse supernovae (SESNe), which have their hydrogen and/or helium layers stripped. We explore several potential progenitor mass diagnostics, primarily the [O I] $\lambda\lambda 6300, 6364$ and the [N II] $\lambda\lambda 6548, 6583$ lines, using their fractional fluxes and time post-peak-explosion as a function for He core mass. These methods have only been tested in models and small sets of supernovae, meaning a more comprehensive approach is needed. By comparing these methods on a large and consistent database of supernovae, we find that the [N II] fitting method produces significantly different results to the more widely used [O I] mass diagnostic, indicating that the [N II] feature may not be a reliable mass diagnostic. We also see that all supernova types (Ib, Ic, IIb) span the entire modeled range of He core masses, which may imply that the stripping of these supernovae are primarily caused by binary companions, rather than being dependent on the properties of the host progenitor stars.

Development of a deterministic adaptive mathematics platform for students transitioning from grade 9 (JHS 3) to grade 10 (SHS 1) in Ghana and Kenya

Ibrahim Adam

Mentors: Franca Hoffmann and Jasmine R. Bryant

This research project studies how an adaptive online mathematics tutor platform can provide unique practice for students transitioning from Grade 9 to Grade 10 in Ghana and Kenya. I have already developed a prototype that organizes mathematical concepts through prerequisite connections, chooses suitable practice questions based on student competency and progress, assesses responses using an online deterministic mathematics assessment system, and uses the resulting evidence to decide subsequent practice questions or advancement for the student. Currently, the system combines verified mathematics questions which covers topics including indices, integer operations, coordinate graphs, ratios and proportions, introductory trigonometry, and similarity. The system supports automatic question rendering, grading and interactive graph-based activities. Some testing that I have done both with the backend on VSCode and the frontend on the live prototype platform has demonstrated successful rendering and grading of both standard and interactive questions, while also

revealing compatibility problems in several older question test cases as well as connecting multi-input interactive activities to the adaptive task logic. Currently, I focused my work on resolving these integration problems, adding more questions to the platform, and making the adaptive progression framework stronger before evaluating learner interaction data.

Toward event-driven feature extraction for neuromorphic brain-machine interfaces

Layla Adeli

Mentors: Azita Emami and Loai Danial

Higher channel counts under lower power budgets and unstable decoding due to electrode degradation have motivated the field of neural interfaces to move towards sparse feature extraction. This project aims to develop a baseline for such an event-driven feature extractor, using an analog front end (AFE) to encode neural activity as pulse-density events. The AFE is used to replace the uniformly-sampled converter and waveform reconstruction that most interfaces require and is the natural input to later feature extraction and decoding. The work was structured in four stages: a comparison of seven unsupervised spike sorters under simulated electrode aging, the construction of a framework for comparing spiking neuron models on the AFE's actual event stream, an evaluation of those neuron models as feature extractors for motor decoding, and the hardware implementation of the logic that could turn the AFE's pulses into a decoding feature. Treating the input as already encoded by the front end, the final stage was implemented in Verilog, including an asynchronous counter of its events. This work contributes toward low-power, on-chip feature extraction for event-driven neural interfaces.

GPU acceleration of TorchLean, a formally verified neural network training library in Lean 4

William H. Adkisson

Mentors: Anima Anandkumar and Robert Joseph George

TorchLean is a library for specifying, training, and formally verifying neural networks. TorchLean is written in Lean 4, an interactive theorem prover that serves as a proof assistant and functional programming language. Because every operation is implemented and checked within Lean, TorchLean avoids the semantic drift that arises when a verified model specification is converted to an external format for verification. This drift can let a proof about the converted model fail to reflect the model that actually runs. This safety came at a steep performance cost, however, since TorchLean's core operations originally ran entirely on CPU, hundreds of times slower than standard machine learning frameworks. This project migrated TorchLean's numerical core to the GPU: first with a custom CUDA matrix multiplication kernel, then with NVIDIA's cuBLAS library, then with targeted use of PyTorch's optimized kernels at the specific operations that profiling identified as dominating runtime. On a roughly 500 million parameter transformer model, these changes reduced TorchLean's slowdown relative to PyTorch from over 700 times to about 4.5 times, while preserving Lean's formal correctness guarantees. Future work will extend formal verification to the GPU operations themselves.

Design and control validation of a five-link planar biped robot

Diya Agarwal

Mentors: Maegan Tucker and Gunter Niemeyer

Bipedal locomotion is a highly compelling topic in robotics due to its relevance to human-like mobility and legged robot design. While progress has been made through complex humanoid bipeds and lower-limb exoskeletons, there remains a need for low-cost, open-source platforms that encourage rapid testing of simple bipedal control strategies in the physical world. This research addresses this through the design and implementation of a low-cost, planar five-link bipedal robot constrained to walk in a single plane using a circular boom mechanism. By constraining motion and ensuring rigid, non-compliant links, the platform provides a clean, controlled environment for studying bipedal walking control. The biped system was also supported with a physics simulation model using MuJoCo to enable accurate real-to-simulation comparisons as well as model and controller validation. Initial experimental results demonstrated the platform's ability to reproduce commanded joint trajectories and achieve walking, providing a foundation for testing and comparing bipedal locomotion controllers on low-cost hardware.

Development of an acoustic bandgap filter for superfluid helium-4 systems

Mitchell J. Agris

Mentor: Keith C. Schwab

Superfluid helium-4 is a promising platform for precision measurements because of its exceptionally low dissipation and long-lived acoustic excitations. However, acoustic energy leakage through fill lines and other supporting structures can significantly reduce the quality factor of superfluid resonators, limiting the performance of a wide range of superfluid systems and circuits. This project investigates the use of periodically crimped tubing as a one-dimensional phononic crystal that suppresses acoustic transmission through Bragg reflection, thereby reducing unwanted energy loss. A custom crimping tool and experimental measurement system were developed to fabricate and characterize periodically crimped waveguides using frequency response measurements. Preliminary measurements on copper and stainless-steel waveguides containing different numbers of periodic crimps revealed clear acoustic bandgaps near the designed Bragg frequency of approximately 12 kHz. The deepest measured bandgap exhibited approximately 60 dB of attenuation relative to an uncrimped waveguide, corresponding to a reduction in transmitted acoustic wave amplitude by roughly a factor of 10^3 and transmitted acoustic power by approximately 10^6 . These results demonstrate that periodically crimped waveguides provide an effective method for suppressing acoustic transmission and establish a promising approach for reducing acoustic loss in future superfluid resonators, circuits, and precision sensing technologies.

High throughput multimodal characterization and correlation analysis of hBN single photon emitters

Sara P. Ahire

Mentors: Harry A. Atwater, Jr., and Kristina L. Malinowski

In this work, we develop a framework for high throughput, multi-modal, single-photon emitter characterization and investigate correlations between coherence-based emitter properties in a large data set. Color centers in hexagonal boron nitride (hBN) generate indistinguishable single photons with high brightness, purity, and coherence. Thus, hBN hosts promising single-photon emitters (SPEs) for scalable quantum photonic technologies, but individual emitter coherence is highly dependent on interactions with the local atomic environment. Therefore, these emitters must be screened for high coherence. Coherence is most directly assessed by resource and time intensive cryogenic, megahertz resolution, linewidth measurements. We seek to identify properties at room temperature that correlate with high coherence at cryogenic temperatures.

We build an automated confocal microscopy platform that acquires photoluminescence spectra and second order autocorrelation ($g^2(\tau)$) measurements for hBN SPEs at room temperature. We correlate the zero-phonon line (ZPL) and full width at half maximum (FWHM) with $g^2(0)$ and other coherence metrics to establish room temperature trends and predict which emitters retain near lifetime limited linewidths at cryogenic temperatures, enabling pre-screening of high coherence hBN sources for quantum photonic applications. Future work will confirm whether the predicted coherence at cryogenic temperatures is accurate.

Composition optimization of an FeCoNiMo high-entropy alloy for hydrogen evolution

Jaehwan Ahn

Mentors: William A. Goddard III and Sejun Kim

High-entropy alloys of earth-abundant transition metals could replace platinum in hydrogen production by water splitting, but their compositional space is too large to search by experiment alone. This project optimized the molybdenum content of an iron-cobalt-nickel-molybdenum alloy on its close-packed (111) surface. Activity was judged by the Volmer step, in which water dissociates and leaves a hydrogen atom bound to the metal. DFT nudged elastic band calculations located the transition state along nine reaction paths. Grouped by the metal atom that receives the hydrogen, the barriers rise from nickel to cobalt to iron. They also follow the local arrangement around each site rather than the overall composition. A MACE interatomic potential fine-tuned on these calculations produced 100

equilibrated surfaces across 10 molybdenum concentrations, and a score built from the barriers ranked them. Activity peaks at 25% molybdenum. Below that value too little molybdenum reaches the surface to form sites, and above it molybdenum atoms sit next to one another and the sites are lost. Work continues on the Heyrovsky step, in which the bound hydrogen leaves as hydrogen gas.

Characterization and optical redesign of the High Contrast Spectroscopy Testbed low-order wavefront sensor

Siria Alicata

Mentors: Dimitri P. Mawet and Clarissa R. Do Ó

Direct imaging of Earth-like exoplanets requires extreme starlight suppression and highly stable optical alignment. This project characterized and improved the long-term pointing stability of the High Contrast Spectroscopy Testbed (HCST) at Caltech, a platform for developing coronagraph technologies for future space missions such as Habitable Worlds Observatory. Extended open- and closed-loop measurements of the Low-Order Wavefront Sensor (LOWFS) showed that active tip-tilt correction reduced residual pointing error by approximately a factor of three. Long-duration experiments also revealed a strong correlation between beam drift and thermal variations, motivating replacement of the original tracking camera and redesign of the LOWFS optical relay. The new system incorporated larger relay optics and a water-cooled Andor Neo camera to improve point-spread-function sampling, reduce vignetting, and increase alignment flexibility. During a 5.3-hour closed-loop validation, the upgraded LOWFS maintained RMS pointing errors below $0.006 \frac{\lambda}{D}$ with negligible long-term guiding drift, providing a more stable platform for future high-contrast coronagraph experiments.

All-optical electron spin coherence in an axial tungsten complex of relevance to tungstoenzymes

Sulaiman Wahdan S W. AlKadi

Mentors: Ryan G. Hadt and Jonathan Aalto

Molecules with unpaired electrons possess angular momentum via spin, an inherently quantum-mechanical degree of freedom. These paramagnetic molecules thus constitute spatially precise, structurally and functionally tunable platforms for quantum sensing. Ultrafast pump-probe optical spectroscopy offers a platform to examine spin dynamics with temporal resolution five orders of magnitude greater than traditional spin initialization and readout methods. However, such spectroscopies require substantive coupling of electron spins to light, which is governed by a complex array of electronic structure factors. In this work, room-temperature, ultrafast all-optical spin coherence is demonstrated in a low-symmetry $S = \frac{1}{2}$ tungsten complex. The electronic structure of this system was probed with absorption, magnetic circular dichroism, resonance Raman, and continuous-wave/pulse electron paramagnetic resonance spectroscopies, as well as quantum-chemical calculations. Ultrafast spin dynamics and contributions to spin decoherence in room-temperature solutions were investigated with time-resolved Faraday rotation. This work both expands the library of optically addressable molecular spin qubits and elucidates contributions to spin decoherence in a system relevant to the active sites of tungstoenzymes, with potential implications for native biological quantum sensing and the electronic design elements required for optical spin-addressability in low-symmetry systems.

Effect of molecular structure on the autoignition behavior of sustainable aviation fuels

Reema B. Almadhi

Mentors: Joseph E. Shepherd and Charline Fouchier

Sustainable aviation fuels can reduce the environmental impact of air transportation, but differences in molecular composition may cause them to ignite differently from conventional jet fuel. This project investigates how molecular features such as branching and cyclic structure influence the autoignition behavior of aviation-relevant fuel mixtures. Experiments are conducted using an automated heated-flask apparatus based on the ASTM-E659 standard method. To validate the automated system, its ability to reproduce the main ignition trends observed in earlier manual and semi-automated experiments was first assessed. The successful validation supported its use for testing the selected fuel blends. Tert-butylcyclohexane and 2,2,4,4-tetramethylpentane were selected and blended with

n-hexane and surrogate paraffinic kerosene. Ignition maps, ignition-probability curves, temperature-variation data, and reaction-delay measurements were used to compare the ignition behavior of each blend with its corresponding reference fuel. This work aims to clarify how molecular structure affects autoignition and to provide guidance for formulating sustainable aviation fuels with ignition characteristics closer to those of conventional aviation fuel.

Tensor compression for faster electron transport calculations

Hadwa I. Alruhaili

Mentors: Marco Bernardi and Tommaso Chiarotti

Electron-electron limited transport properties are important for understanding collective electronic behavior in quantum materials. In particular, first-principles calculations of quantities such as electronic viscosity require accurate electron-electron scattering rates. These calculations are currently limited by the cost of evaluating and storing the screened electron-electron interaction, which is represented by a high-dimensional tensor that depends on several momentum and band indices. Compression methods have already been applied successfully to other large quantities in transport calculations, including electron-phonon matrix elements, where tensor-based approaches have reduced storage and interpolation costs. Here, we aim to apply a similar tensor compression strategy to the screened electron-electron interaction. The interaction data are stored in an HDF5 file with separate real and imaginary components, and a PyTorch data-loading pipeline has been developed to read small tensor blocks without loading the full dataset into memory. A preliminary low-rank model was also tested to verify the full workflow, including batching, GPU transfer, prediction, loss evaluation, and back-propagation. Although this initial model produced poor reconstruction accuracy, it confirmed that the computational pipeline is working. The next step is to implement a physics-consistent tensor hypercontraction model. The final goal is to achieve at least an order-of-magnitude reduction in memory while maintaining the accuracy of the original first-principles interaction tensor.

Census of low-mass galaxies with the Hubble Space Telescope

Abrar Amin

Mentor: Mike G. Jones

Dwarf satellite galaxies, small galaxies gravitationally bound to a larger host, are valuable for studying galaxy formation and dark matter because of their high mass-to-light ratios. Most of what we know about satellite systems comes from the Milky Way's own, and whether it is representative of Milky Way-like galaxies in general is an open question. Answering it requires knowing which nearby dwarf galaxies are truly bound to their hosts rather than lying behind them, which in turn requires precise distances. We measure distances using the tip of the red giant branch, the predictable peak brightness low-mass stars reach at the end of the red giant phase, in dwarf galaxy candidates imaged by the Hubble Census of Nearby Satellites. From each galaxy's stellar photometry, we select red giant branch stars by color, brightness, and signal quality, after correcting for dust in our own galaxy. We then fit the tip brightness by maximum likelihood, comparing the observed distribution of stellar brightnesses to a model that accounts for measurement error and for the fraction of stars recovered at each brightness. These distances separate true satellites from background galaxies across the survey, placing the Milky Way's satellite system in context among its peers.

Identifying and characterizing dipper stars in Taurus

Anica J. Ancheta

Mentors: Eric J. Gaidos, Andrew W. Howard, and Gabriella Zsidi

Dipper stars are a subclass of young stellar objects that exhibit frequent dimming events caused by the dust in the circumstellar disk obscuring the observer's view of the stars' light. The processes that cause dipper behavior in these stars lift the dust above midplane in the innermost part of the disk, enabling us to characterize dust properties where terrestrial planets are expected to form. In this work, we analyze a multifilter photometric dataset from Las Cumbres Observatory's 0.4-meter network of telescopes. The sample of dipper stars comes from the nearby Taurus star-forming region,

which hosts a $\lesssim 3$ Myr population of young stars. Here, we determine the stars exhibit dipping behavior during the observing seasons in the Taurus region, then model the dust grain size distributions of each star. We then look for correlations with stellar and disk parameters to identify whether dipper behavior is linked to specific physical conditions.

A multi-wavelength archival search for counterparts and hosts of *NuSTAR* fast X-ray transients

Tooba Ansar

Mentors: Fiona A. Harrison and Guillem Megias Homar

Fast X-ray transients (FXTs) are short-lived bursts of X-rays whose origins remain poorly understood. While recent advancements in soft X-ray surveys have enabled rapid follow-up on and identification of multi-wavelength counterparts and host galaxies to several soft X-ray FXTs ($E < 10$ keV), the nature of FXTs identified in archival hard X-ray observations is largely unexplored. *NuSTAR*, which observes the hard X-ray sky in the 3–79 keV range, provides a unique opportunity to investigate these transients at higher energies. In this work, we conduct a systematic multi-wavelength search for counterparts and host galaxies of 52 FXTs found in archival *NuSTAR* data. Searches of optical, radio, and ultraviolet survey data reveal no convincing contemporaneous transient counterparts. However, cross-matching with *Gaia* within the *NuSTAR* localization radius reveals 17 FXTs with nearby Galactic stellar sources, including two potential M-dwarf candidates that could indicate Galactic origins. We also identify two promising candidate host galaxies with probabilities of chance coincidence $P_{\text{chance}} \lesssim 0.1$. Our results highlight the limitations of archival searches and the lack of contemporaneous follow-up in identifying promising counterparts, and emphasize the importance of prompt, coordinated follow-up to future FXTs in constraining the origins of this enigmatic transient population.

Alkyne synthesis by the cross-coupling of nitroalkanes with nitroalkenes

Adais M. Arora

Mentors: Marvin Parasram and Brian M. Stoltz

The alkyne is a versatile functional group with wide applications ranging from natural products to medicinal compounds. Compared to the numerous strategies employing alkynes, their synthesis is relatively underdeveloped. The Parasram lab aims to develop a mild, modular, and user-friendly alkyne synthesis. This method is a transition-metal free cross-coupling of nitroalkanes with nitroalkenes that will serve as an effective complement to existing alkyne synthesis methods. My work specifically focuses on demonstrating the scope of this method by first synthesizing structurally interesting nitroalkene substrates and then attempting to convert them into alkynes via the Parasram lab's method. Substrates were synthesized successfully through a Henry reaction with nitromethane and selected aldehydes. Isolated substrates have been obtained after purification through silica column chromatography with yields ranging from ten to thirty percent.

Variability in broad emission line components of potential intermediate-mass black holes

Kiera E. Aston

Mentors: Grace M. Olivier and Michael E. Brown

Supermassive black hole (SMBH) formation could greatly impact galactic growth and formation since a SMBH lays at the center of each large galaxy we know in the local universe. Studying IMBH behavior could reveal much about SMBH formation and thus galactic growth and formation. Blue compact dwarf galaxies (BCDs) are not massive enough to host SMBHs, so if they have active galactic nuclei (AGN) they would be intermediate-mass black holes (IMBHs). However, relatively few IMBH candidates have been identified. We are looking for variability in the broad components of visible spectroscopic emission lines to identify IMBH candidates. Broad emission components are generated by high velocity gas creating a doppler effect. Looking primarily at BCDs with previously identified broad components, we will examine how they have changed over time. We will compare the flux percentage from data taken using the Magellan Echellette (MagE) Spectrograph in spring 2026 to percentages reported in

2007. Variability was identified in one of the five galaxies in our sample by comparing the percent of flux from the broad component of the H α line of J1143 over time. A significant change was found, indicating J1143 is a likely host of an IMBH. Identifying variability in the broad components of other galaxies would indicate a likely IMBH candidate.

Enantioselective crystallization of an RNA precursor and the origin of homochirality

Carolina Ayala-Garnica

Mentors: S. Furkan Ozturk and Kyogo Nagashima

Ribose-aminooxazoline (RAO) is a key intermediate in prebiotic ribonucleotide synthesis and is hypothesized as the origin of RNA's homochirality. RAO forms conglomerate crystals, a rare property where enantiomerically pure crystals individually form from racemic mixtures. Recent research demonstrates that RAO enantioselectively crystallizes on a magnetic, mineral surface. During enantioselective D-RAO crystallization on a magnetite surface, L-RAO remains in solution and can crystallize in a nearby environment resulting in racemic mass balance in the environment. This diminishes the enantioselectivity of the process on a macroscopic scale and makes it difficult to maintain homochirality. We offer a solution through phosphate-catalyzed hydrolysis of L-RAO in solution to ribose oxazolidinone, leaving homochiral D-RAO crystals intact. By hydrolyzing L-RAO in solution, we break the mass balance between enantiomers and maintain D-RAO homochirality for subsequent ribonucleotide synthesis.

Shaping bio-realistic spiking neural network dynamics via gradient-based optimization of connectivity parameters

Maggie Y. Bai

Mentors: Costas Anastassiou and Shayan Farzad

Neural oscillations are a hallmark of coordinated brain activity, essential for communication, sensory processing, and cognition, and their disruption is implicated in disorders such as schizophrenia and epilepsy. While biophysically detailed models can faithfully capture the behavior of individual neurons, the circuit-level parameters that govern network dynamics--synaptic strengths, connection probabilities, and time constants--are far less constrained by data, and their mapping to emergent activity is highly nonlinear. Tuning these parameters to produce a target behavior therefore typically relies on slow, gradient-free search that can require many times more simulations than gradient-based methods. This project investigates whether Jaxley, a differentiable simulator that computes exact gradients through networks of morphologically detailed neurons, can instead shape emergent dynamics directly by optimizing only connectivity parameters. Working with recurrent excitatory-inhibitory networks that produce gamma-band oscillations, we assess how well gradient-based optimization can steer a circuit toward a defined oscillatory state and benchmark its efficiency and reliability against existing approaches--aiming to make large-scale, biologically realistic brain simulation more tractable.

Diffusion guided posterior sampling for Bayesian inverse problems

Nicholas Z. Bai

Mentors: Anima Anandkumar, Jiachen Yao, and Xi Deng

Bayesian inverse problems seek to infer unknown physical fields or parameters from indirect, often sparse observations through a forward model—for example, reconstructing a permeability field from pressure measurements governed by a PDE. Because these problems are typically ill-posed, the solution is characterized by a posterior distribution rather than a single maximum-a-posteriori estimate. Existing sampling methods can recover locally plausible solutions but often fail to capture the full posterior when it is strongly multimodal. We address this challenge by separating posterior sampling into two complementary objectives: correctly allocating probability mass across modes and accurately resolving the distribution within each mode. Through theoretical analysis and controlled experiments, we find that gradient-based particle transport is effective for local refinement but poorly

suited to reallocating mass between well-separated modes. Motivated by this observation, we develop a hybrid framework that combines mode-level resampling for global mass allocation with particle transport for local refinement. We further investigate how learned generative priors can accelerate inference while maintaining reliable posterior sampling.

Developing a computational optimization framework for cell-class-specific electrical stimulation

Ishita Banerjee

Mentors: Costas Anastassiou and Philip Wong

Electrical stimulation (ES) is widely used to modulate neuronal activity in both basic neuroscience research and clinical neuromodulation therapies. However, designing stimulation protocols that selectively target specific neuronal populations remains a significant challenge due to the complex relationship between network parameters and neuronal dynamics. This project aims to develop a computational optimization framework for identifying network parameters that reproduce target electrophysiological behavior in biophysically detailed neuronal networks. To investigate this problem, a recurrent network composed of six parvalbumin (PVALB) interneurons was constructed in BMTK, and a randomized parameter search was implemented to optimize excitatory synaptic weight, inhibitory synaptic weight, synaptic delay, and stimulus amplitude using a loss function based on voltage trace similarity and firing rate. Optimization performance was evaluated across twenty independent trials through analyses of voltage traces, firing rates, optimization loss, and parameter recovery. The optimization framework successfully identified parameter sets that generated network dynamics similar to those of a reference network, with the lowest-loss solutions closely reproducing the target voltage traces and firing behavior. The results also demonstrate that multiple parameter combinations can produce comparable network activity, highlighting the challenges of parameter recovery within high-dimensional optimization spaces. These findings establish an initial computational framework for network-level parameter optimization and provide a foundation for developing more efficient optimization strategies for computational models of electrical stimulation.

Understanding the cataclysmic variable population in the globular cluster 47 Tucanae

Avni Bansal

Mentors: Kareem J. El-Badry and Geoffrey Mo

We search for new cataclysmic variables (CVs) and characterize previously known CVs in the globular cluster 47 Tucanae using ~400 days of HST WFC3/UVIS F300X imaging. We identify several new CV candidates on the basis of variability, CMD position, and X-ray brightness. We also measure a period of 173.7 minutes for the suspected magnetic CV W62, providing evidence against a magnetic classification. Finally, we observe dwarf nova outbursts in two sources, W25 and W36. Our data provides the first well-sampled coverage of an outburst in W25 and the first coverage of the DN decline in W36. We model both events and interpret them in the context of various statistical relations between DN features, as well as the thermal-viscous accretion disk instability model. We also discuss how W36 departs from standard CV behaviour because its companion is a subgiant rather than a main sequence star.

Characterization of the organ systems and vasculature of sleeping and exercised glassfrogs

Anna S. Becker

Mentors: Carlos A. Taboada and Patrick Simpson

Glassfrogs are a type of frog that turn transparent while sleeping as a form of camouflage. Previous work has found that glassfrogs enter this state by storing ~89% of their red blood cells (RBCs) in their liver. However, the anatomical changes in glassfrogs that occur during camouflage are not yet documented. Furthermore, it is unknown how the flow of oxygenated blood could change to account for a lower concentration of RBCs in circulation during camouflage. Recent work has found that after blue-light induced release of stored RBCs, the anterior abdominal vein (ABV) appears oxygenated, which indicates that a branch of this vein connects directly to the heart. To determine the ABV's branches and the organ anatomy, we analyzed MicroCT images of both the sleeping and exercised glassfrogs, and created organ and blood vessel models with 3D slicer. We found significant changes in

structure, size and positioning in the frog's kidneys, lungs, stomach and liver. Additionally, we tracked the ABV through the liver and found possible connections between the heart and the vein, which suggest an alternative path for circulation during camouflage. These models and observed changes of the glassfrogs' organs and vasculature provide the basis for further exploration of glassfrogs' mechanism of camouflage.

Laser-grown superconductors: High-Tc niobium thin films via thermal laser epitaxy

Sophia M. Benedick

Mentors: Austin J. Minnich, David S. Catherall, and Yifei Yan

Niobium (Nb) is a type-II superconductor with the highest elemental T_c of 9.2 K in bulk. Nb thin films are of interest, as they are used to create Josephson junctions for quantum computing and superconducting quantum interference devices (SQUIDs). However, because Nb has a high melting point of 2477 C, typical thin film growth techniques, such as molecular beam epitaxy (MBE) and e-beam evaporation, present challenges, including effusion cell failures and slow deposition rates. Therefore, we turn to thermal laser epitaxy. This physical vapor deposition technique directs a continuous-wave 1 kW fiber laser (1070 nm) onto a 3 mm diameter Nb target mounted inside an ultrahigh-vacuum chamber. The laser vaporizes the Nb, which then condenses on a 1 x 1 cm² c-plane sapphire substrate 70 mm above it. With a T_c of 9.22 K and a residual resistance ratio of 5.82, our 25 nm Nb thin film is of excellent quality. This work demonstrates the use of thermal laser epitaxy in growing high-quality high-Tc superconducting thin films, extending its application to thin film deposition of in demand materials.

Molecular basis of CK1δ nuclear import in circadian rhythm modulation

Emma J. Berg

Mentors: André Hoelz and Chia-Yu Chien

The circadian clock is an internal biological timing system that aligns physiological processes with external light-dark cues. At the molecular level, it is regulated by a transcription-translation feedback loop involving the clock proteins PER and CRY, whose activities are modulated by Casein Kinase 1δ (CK1δ). CK1δ is reported to be localized both in cytosolic and nuclear compartments, and its transport between the two is mediated by the karyopherin (KAP) family. However, the detailed molecular mechanisms governing the shuttling of CK1δ and how it participates in distinct stages of signaling remain elusive. This project aims to validate the predicted nuclear localization signal (NLS) of CK1δ and determine how its interactions affect its binding to α-KAPs and how it enters the nucleus. To accomplish this, we tested the interaction between CK1δ and mKapa2 through Size-Exclusion Chromatography coupled with Multi-Angle Light Scattering (SEC-MALS). The results suggested that CK1δ binds α-KAPs with low affinity, and the C-terminal extension might be essential for the interaction. Additionally, we conducted subcellular localization and Western blot analysis to visualize CK1δ's localization within mammalian cells. Aligned with our biochemical findings, we found that CK1δ localizes to both the nucleus and cytosol, which demonstrates a property of low-affinity NLS-tagged cargoes.

Safety and contingency certificates for drone landing

Juan E. Beron

Mentors: Aaron D. Ames, Joe Moeller, Pedro Roque, and Yana Lishkova

Autonomous landing in cluttered and changing environments requires more than reaching a designated target: the vehicle must remain safe while preserving the ability to divert when conditions deteriorate. This work develops a safety-critical autonomy architecture that separates nominal task execution from safety and contingency enforcement. Model predictive control generates dynamically feasible nominal commands, while Poisson safety functions transform complex obstacle geometry into smooth differentiable fields that are incorporated into control barrier function constraints for collision avoidance. Control Lyapunov function-based regions characterize convergence and target reachability, and an r-out-of-p combinatorial condition preserves a prescribed number of feasible alternatives throughout the maneuver. Ellipsoidal funnel constructions are further investigated as a mechanism for maintaining connectivity through cluttered spaces. The framework is evaluated through simulation and

Crazyflie hardware experiments with OptiTrack feedback, demonstrating real-time integration of predictive control, safety filtering, and contingency-aware navigation. Together, these components provide a modular foundation for autonomous systems that must pursue a primary objective without prematurely sacrificing safe alternatives under geometric constraints and environmental change.

SNAP: An SSL-based neural network for automated pseudo-time generation

Rishi Bhargava

Mentors: Matthew W. Thomson and Hao Wang

Infertility affects one in six people worldwide, yet fertility clinics cannot reliably predict which embryos will develop successfully for in-vitro fertilization (IVF). Embryo viability is assessed using preimplantation genetic testing, which requires an invasive biopsy and can be confounded by mosaicism, or by morphological grading, which is subjective and shows low inter-observer agreement for some features. Computer vision offers a non-invasive and objective alternative by extracting quantitative developmental information directly from embryo images. Our long-term aim is to read developmental potential from a single image. Here, we address a critical first step: determining whether developmental time itself can be recovered from static embryo morphology. Cleavage timing correlates with embryo viability and, unlike viability, can be directly measured from time-lapse recordings, providing an objective ground truth for model training. We therefore ask whether a single label-free image can reveal where a zygote lies within its first cell cycle and how accurately this information can be recovered relative to existing approaches applied to identical data. We find that a single image can classify an embryo's cleavage timing to within 1.3 hours across the roughly 18-hour span of the first cell cycle, and does so more accurately than all seventeen alternatives we scored on the same images, including two published embryo staging schemes and three frontier general-purpose large language models.

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.

Topological phase transitions in three-dimensional quantum spin liquids with neural quantum states

Sanzhar Bissenali

Mentors: Norman Yao, Marco Bernardi, and Sajant Anand

Neural quantum states (NQS) offer a variational route to strongly correlated phases of matter but optimizing them in topologically ordered regimes remains difficult. We extend the approximately symmetric neural network ansatz, previously developed in two dimensions, to the three-dimensional toric code in a uniform magnetic field. While symmetry-unaware convolutional networks get stuck in local minima, our ansatz converges to quantum Monte Carlo (QMC) ground-state energies with relative errors of order 10^{-4} - 10^{-3} for systems of up to 882 spins. Using the Fredenhagen-Marcus string

order parameter and the Rényi entropy, we map the topological-to-trivial phase diagram in the sign-problem-free (h_x, h_z) plane and find good agreement with prior QMC predictions. The same architecture applies directly to the sign-problem-full regimes (the h_y field and the fermionic toric code) where unbiased QMC is unavailable.

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 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.

A Search for FU Orionis-like objects in the APOGEE and BOSS Young Star Survey

Valentina S. Bonilla Villalobos

Mentor: Lynne Hillenbrand

FU Orionis objects (FUors) represent the most luminous and long-lasting class of outbursts from young stellar objects, driven by episodes of dramatically enhanced accretion. To date, only a few dozen FUors have been identified, primarily through photometric monitoring of their characteristic light-curve behavior. However, the increase in the discovery rate of FUors since the establishment of time-domain surveys suggests that many more FUors may remain undiscovered in their fading phase. In this work, we investigate the potential for spectroscopic identification of FUors within the APOGEE and BOSS Young Star Survey (ABYSS), an SDSS-V program acquiring optical and near-infrared spectra for approximately 100,000 young stellar objects. We examine the ABYSS target samples and selection criteria to assess the likelihood of including FUor candidates. Using high-resolution H-band IGRINS spectra of known FUors, we identify prominent spectral features and measure their equivalent widths. These measurements are compared with those of ABYSS APOGEE sources to define a diagnostic parameter space characteristic of FUors and to identify high-confidence FUor candidates within the survey. Our results demonstrate the potential of large spectroscopic surveys to significantly expand the known FUor population and provide new insights into episodic accretion in young stars.

Creating AI-ready multi-modal datasets for cancer research

Aditya P. Borate

Mentors: Ashish Mahabal and Sean Kelly

The Early Detection Research Network holds multi-modal cancer collections that are scientifically valuable, yet difficult to reuse at scale because cross-file relationships, variable semantics, and acquisition context are often embedded in naming conventions, spreadsheets, and publications rather than machine-actionable schemas. This project evaluates and improves the readiness of selected collections for artificial intelligence using a ten-dimensional scoring framework and standardized dataset descriptors based on the MLCommons Croissant format. We completed seven assessments across flow cytometry, mammography, liquid chromatography mass spectrometry, and protein microarray, with readiness scores on a 0 to 100 scale ranging from about 43 to 81. Imaging ranked

highest; proteomics and microarray deposits were strong scientifically, but weaker in operational packaging. We mapped dependencies among readiness dimensions and found that data quality and structure scored well while operational readiness scored lowest across collections, indicating that gaps lie in documentation and reusable packaging rather than scientific content. Parallel work on cell segmentation and pancreas segmentation showed the same pattern when acquisition context was not harmonized. These collections are scientifically strong but often need modest metadata and packaging improvements. We will extend Croissant descriptors to additional datasets and assess the broader network of collections.

Investigating the interactions of the NEIL3 beta hairpin and KY motif with MCM2–7 in DNA interstrand cross-link repair

Hannah L. Bowman

Mentors: Daniel R. Semlow and Gabriel Ashton-Rickardt

DNA interstrand cross-links (ICLs) covalently link the two strands of DNA, stalling replication forks and threatening genome stability. One class of ICLs, abasic site-induced ICLs (AP-ICLs), is repaired by the NEIL3 glycosylase. Although several domains within the NEIL3 C-terminal domain have established roles in replication-coupled ICL repair, the functions of two highly conserved motifs – the KY motif and the β -hairpin – remain poorly understood. Our lab's AlphaFold modeling suggests that the β -hairpin interacts with a cleft in the MCM2–7 helicase, potentially positioning the glycosylase domain of NEIL3 to unhook ICLs at stalled replication forks. To investigate this model, we generated a panel of NEIL3 mutants containing deletions, alanine substitutions, and targeted point mutations within the KY motif and β -hairpin to test the contribution of residues predicted to mediate interactions with MCM2–7. Following DNA cloning, baculovirus-mediated expression in Sf9 insect cells, and protein purification, mutant proteins will be evaluated in *Xenopus* egg extract using biochemical rescue-depletion assays to determine whether these motifs are required for NEIL3-dependent ICL repair. These studies will clarify how the KY motif and β -hairpin support replication-coupled ICL repair and provide a foundation for future studies examining how these motifs mediate interactions with the MCM2–7 helicase.

Disturbance-aware trajectory planning with learned linear-in-latent disturbance dynamics

Erik S. Bryland

Mentors: Soon-Jo Chung and Min Kim

High-precision autonomous maneuvers may fail in the presence of onboard disturbance sources, such as a sloshing liquid payload. In this study, disturbances are predicted as outputs of a latent state-space model whose time-varying coefficients are learned functions of observable quantities. Using this model, we implement an open-loop Monte Carlo Tree Search trajectory planner that incorporates the learned disturbance model. The planner's action sets are generated with spectral analysis. Chance constraints are incorporated through Monte Carlo sampling. We validate the planner on two nonlinear systems with hidden disturbance dynamics, including a system of one-dimensional coupled Duffing oscillators and a two-dimensional system of particles coupled by a nonlinear spring. The performance of the learned model is evaluated against a Koopman-based approach across two tasks: disturbance prediction and downstream planner performance. The planner demonstrates successful disturbance-aware trajectory generation by rerouting around obstacles while accounting for state uncertainty.

Actor-critic reinforcement learning for model predictive control for the Autonomous Indy Car Challenge

Ettienne Caban-Klepac

Mentors: Soon-Jo Chung and Gilchrist X. Johnson

This project implements AC4MPC (Reiter et al., arXiv:2406.03995), in which a reinforcement-learning agent is trained on the exact cost a model predictive controller (MPC) optimizes, and its learned value function (critic) is embedded in the MPC as a terminal cost. The terminal cost summarizes the infinite-horizon consequences of a state, granting the finite-horizon controller long-range foresight, while the trained policy warm-starts the solver. The correctness requirement is that the RL reward equals the negative of the MPC's own stage cost; only then is the critic a valid terminal value function for this controller. To guarantee this, we exposed the production single-track racing MPC's dynamics, stage

cost, and 23 of 28 constraint rows to Python via pybind11, so the environment steps the exact dynamics and is scored by the exact cost. A Soft Actor–Critic agent with smooth, compact networks was trained, exported to ONNX with a verified value-scale inversion (relative error 10^{-7}), and integrated as a first-order terminal-cost term in the C++ solver. The full pipeline runs in production

on Laguna Seca with zero failed solves. Experiments show the critic faithfully transmits its learned preference, driving closed-loop behaviour across the full range from over-conservative (3.2 m/s) to over-aggressive (solver-breaking), and can be tuned to a stable operating point near the 17.7 m/s nominal baseline. Achieving a net speedup is now a training-scale question rather than an integration one: single-step SAC at 2×10^5 steps has not yet learned feasible high-speed cornering, versus the reference paper's 2×10^6 . the negative of the MPC's own stage cost; only then is the critic a valid terminal value function for this controller. To guarantee this, we exposed the production single-track racing MPC's dynamics, stage cost, and 23 of 28 constraint rows to Python via pybind11, so the environment steps the exact dynamics and is scored by the exact cost. A Soft Actor–Critic agent with smooth, compact networks was trained, exported to ONNX with a verified value-scale inversion (relative error 10^{-7}), and integrated as a first-order terminal-cost term in the C++ solver. The full pipeline runs in production on Laguna Seca with zero failed solves. Experiments show the critic faithfully transmits its learned preference, driving closed-loop behaviour across the full range from over-conservative (3.2 m/s) to over-aggressive (solver-breaking), and can be tuned to a stable operating point near the 17.7 m/s nominal baseline. Achieving a net speedup is now a training-scale question rather than an integration one: single-step SAC at 2×10^5 steps has not yet learned feasible high-speed cornering, versus the reference paper's 2×10^6 .

Assessing student engagement across Math Camp activities in Kenya via a mixed-methods evaluation of activity design

Giovanni Cabrera

Mentors: Franca Hoffmann and Jasmine R. Bryant

The math camps initiative aims to pique interest in mathematics through short, activity-based sessions. The specific features of activity design that drive student engagement has not been evaluated, particularly in the African context. This project assessed student engagement across eight math camp activities delivered in Kenya, combining a quantitative survey with a qualitative review of open-ended student responses along with facilitator lesson plans. Immediately after each activity, students completed a short survey rating their engagement on a seven-item scale, along with two open-ended questions. Surveys were collected on paper and later transcribed using an artificial-intelligence-assisted process, checked against the original documents. Facilitator lesson plans were reviewed and six design features associated with effective discovery-based learning were extracted from them. Engagement was consistently moderate to high across all eight activities (composite scores of 3.5-4.1 out of 5), and nearly every activity shared the same core design features, particularly guided discovery and productive struggle, which limited the ability to isolate any single feature's effect using this dataset alone. Open-ended responses, however, provided independent qualitative evidence that students engaged meaningfully with each activity and their target mathematical idea. Students often referenced concepts such as parity, powers of two, and prime numbers. These findings suggest the camp's activities succeeded at their stated goal of building interest in mathematics.

Morphological and immune response to gas vesicle cavitation by focused ultrasound in mouse glioma models

Michelle Cai

Mentors: Mikhail G. Shapiro and George H. Daglian

Abstract withheld from publication at mentor's request.

Soft ionic polymers for stretchable ionic circuits

Can S. Canpolat

Mentors: William A. Goddard III, Chiara Daraio, Max Tepermeister, and Sejun Kim

Wearable bioelectronics need thin-film circuits that stretch with the body, which requires an elastic, ionically conductive material between their rigid components. In this project, we develop a non-swelling, non-brittle, tough, elastic, and stretchable hydrogel for a glass capillary setup used to probe the electrical properties of the ionic circuit. An initial polyacrylamide hydrogel proved too brittle under tensile testing. We later evaluated four alternative formulations: a poly(vinyl alcohol)/zwitterionic monomer double network, a copolymer of a zwitterionic monomer with acrylamide, and a concentrated, single-monomer zwitterionic gel. The first two failed through spontaneous precipitation and excessive swelling. The zwitterionic copolymer polymerized only partially, suggesting that the initiator needs to be changed. We attribute the incomplete cure to a photoinitiator different from the one in its literature basis. We predict that the concentrated zwitterionic gel offers the most promising result with the highest toughness and greatest elasticity while preserving swelling. Remaining work includes repeating that formulation with the intended photoinitiator, switching to a redox initiator for the copolymer route, and falling back to two other monomer combinations if neither succeeds.

Photonic molecule design for quantum frequency conversion

Matthew J. Cantor

Mentors: Kiyoul Yang and Kerry J. Vahala

Scalable quantum networks require efficient interfaces between various quantum processors that emit and absorb photons at different optical frequencies. This project develops an integrated quantum frequency converter linking two visible-band quantum emitters separated by a frequency shift of tens of terahertz. The device utilizes Bragg-scattering four-wave mixing in a silicon nitride photonic molecule composed of two coupled high-Q microring resonators. Two optical pumps mediate the conversion and hybridize across the coupled resonators, providing a tuning degree of freedom that satisfies energy conservation while keeping all four waves aligned to cavity resonances. The signal and idler, however, are localized in a single ring to maintain strong nonlinear overlap for high conversion efficiency at reduced pump powers. The design methodology combines electromagnetic mode simulations with coupled-mode theory and has identified geometries that satisfy the conditions of Bragg-scattering four-wave mixing while suppressing competing nonlinear processes. The next steps are to fabricate the device and characterize its performance. This work establishes a pathway toward low-noise, chip-scale quantum interconnects capable of linking dissimilar quantum-network nodes.

Constructing a SNIPR-based genetic circuit for non-invasive acoustic imaging of glioblastoma-specific antigens

Barbara S. Cao

Mentors: Mikhail G. Shapiro and Ann Liu

Abstract withheld from publication at mentor's request.

Control of surface morphology in spalled thin silicon solar cells for space solar power applications

Enya Cao

Mentors: Harry A. Atwater, Jr., Andy Nyholm, and Elsie P. Loukiantchenko

Space solar power requires photovoltaic systems with high power-to-mass ratios, motivating thin silicon devices. Spalling offers a scalable, kerfless route to exfoliate thin layers of silicon substrate. Since Si is an indirect bandgap material, thinning reduces optical absorption of photons. However, controlled corrugations may enhance light trapping due to increased scattering. The effect of processing conditions on spall morphology remains poorly understood. We investigate the effects of wafer orientation, crack-propagation direction, and electroplating conditions using surface characterization tools and finite element analysis. Si(100) commonly develops crystallographic corrugations, while river lines occur in both Si(100) and Si(111) and intensify near lateral edges. A 3D FEA COMSOL model is built to connect this surface feature to an opposite-signed, edge-enhanced out-of-plane shear and mixed-mode I+III fracture. Directional studies indicate that Si(111) spalling along the $\langle 110 \rangle$ orientations generally produces smoother surfaces than along the $\langle 211 \rangle$ orientations, while Si(100) results primarily reveal smooth, corrugation-free regions along the $\langle 100 \rangle$ orientations. Electroplating conditions remain ongoing work and are expected to reveal how stressor residual stress gives rise to distinct morphological features.

Roll-slip detection for robotic tactile sensing in contact-rich manipulation

Susanna S. Cao

Mentors: Monroe D. Kennedy, Joel W. Burdick, and Jinho So

Tactile sensing enables humans to navigate their environment by identifying object geometry, material properties, and contact interactions without visual feedback. While current research provides methods for manipulative object and material identification, little work has been done on detecting force interactions in contact-rich manipulation tasks. Current methods typically rely on a single 6-axis measurement, which provides only a net resultant force and torque. This project explores whether high-resolution robotic tactile sensing can be used to detect rolling and slipping to better understand contact interactions between objects. To do so, we first constructed 2mm resolution piezoelectric tactile sensors with a silicone layer to improve shear force sensitivity. Tactile data are then collected to compare the rotational displacement of the lid to the robotic hand trajectory to distinguish rolling from slipping. A neural network is then constructed to correlate temporal tactile measurements to a roll-slip ratio to estimate the normal and frictional contact forces between the objects to guide adjustments and alignments. This research aims to establish a framework for determining the roll-slip behavior using tactile sensing for applications such as threaded lid assembly or O-ring installation for prosthetic arms and dexterous manipulation.

PTNO: Training neural operators with noisy Monte Carlo estimates for particle transport problems

Yubo Cao

Mentors: Anima Anandkumar and Xi Deng

Particle transport under multiple scattering is central to radiative transfer and plasma physics, yet high-fidelity Monte Carlo (MC) simulations demand tracing prohibitively large numbers of particles. Learning-based surrogates offer a path to amortize this cost, but their training typically requires expensive, well-converged MC solutions. We propose Particle Transport Neural Operator (PTNO), a neural operator framework for learning particle transport surrogates directly from noisy, low-cost MC labels. Training on such labels presents two coupled challenges: (1) the MC labels have high variance, which destabilizes standard supervised learning, and (2) particle transport solutions exhibit high dynamic range (HDR), spanning many orders of magnitude and making the solution difficult for a surrogate to resolve. To address the first issue, we train on noisy labels of many different configurations to learn the solution operator of the PDE, amortizing the cost of MC runs and enabling zero-shot generalization. For the second issue, a standard approach is to apply a non-linear transformation, such as log, to compress the HDR. However, this makes the resulting supervision biased. Instead, PTNO learns in the original space to keep the labels unbiased by adding an output

softplus layer to PTNO, enforcing positivity. This allows small positive values to be represented smoothly and avoids nonlinear transformations of noisy MC labels. We demonstrate PTNO's effectiveness on representative particle transport problems, including neutron transport in fusion-reactor settings and radiative transfer in spatially varying participating media, showing that neural operators can learn accurate cross-domain transport surrogates while delivering 10^5 -- $10^6\times$ faster inference compared to converged MC simulations.

The degree of the symplectic Stiefel manifold

Ezra Caspary

Mentor: Yassine El Maazouz

The symplectic Stiefel manifold is the set of linear maps preserving the canonical symplectic form. It has applications in physics, engineering, and robotics, specifically in the reduction of Hamiltonian systems and optimization under symplectic constraints. The symplectic Stiefel manifold can be realized as an algebraic variety cut out by quadratic relations arising from matrix equations which enforce the symplectic condition. The degree of the variety, a fundamental geometric invariant, is unknown and extremely expensive to compute symbolically even for individual cases. We use tools from commutative algebra, algebraic geometry, and representation theory to determine the degree of the symplectic Stiefel manifold in general.

Spectral methods for causal effect estimation in linear time-series models

Rafael I. Castro

Mentors: Bijan H. Mazaheri and Leonard J. Schulman

We study causal inference in linear time-series systems and perform soft interventions on the causal variables, driving them at chosen frequencies and taking the Fourier projection of observed signals to separate causal structure from confounding. When estimating the effect α of one variable X on another Y and using a soft driving intervention $D(t) = e^{i\omega^*t}$ with effect strength δ , we find the signal-to-noise ratio is proportional to the quantity $\delta\omega^*\sqrt{T}$, increasing with stronger, higher frequency, and longer drives. We then extend the mechanism to DAG recovery on n nodes by driving each node simultaneously at a distinct frequency and de-mixing the response from a single time record. This scheme improved the error of each recovered edge by a factor of $\sqrt{n}$ for equal total observation time. Our results were verified on simulated systems for random Erdős-Rényi DAGs.

Automated body composition analysis for cardiac risk prediction in Duchenne muscular dystrophy

Christian W. Cavener

Mentors: Shreyas S. Vasanawala and Konstantin Zuev

Duchenne muscular dystrophy (DMD) is an X-linked progressive neuromuscular disorder in which cardiomyopathy remains a leading cause of death. Skeletal muscle atrophy and fat deposition visible on magnetic resonance imaging (MRI) may provide quantitative biomarkers for predicting cardiac decline and evaluating treatment response in patients. However, MRI-based body composition analysis in patients with DMD is still largely performed manually or semi-automatically, limiting its scalability and routine clinical use. This project aims to develop an artificial intelligence-based pipeline for automated tissue segmentation and body composition analysis from cardiac MRI scans in patients with DMD. A retrospective cohort of approximately 200 patients with DMD undergoing routine cardiac MRI will be identified. Axial localizer images will be manually segmented using 3D Slicer to generate ground-truth labels for subcutaneous adipose tissue, visceral adipose tissue, epicardial adipose tissue, intramuscular adipose tissue, and skeletal muscle. Skeletal muscle fat fraction will also be calculated. A U-Net-based convolutional neural network will be developed using Python, PyTorch, and MONAI and will utilize transfer learning from existing tissue segmentation frameworks such as nnU-Net. Both volumetric segmentation across selected vertebral levels and single-slice segmentation at a standardized mid-thoracic level will be evaluated. Model performance will be evaluated by comparing automated predictions with manual tissue segmentations using Dice similarity coefficients and Jaccard indices. The resulting quantitative body composition measurements may ultimately support improved prediction of cardiovascular outcomes and treatment response in patients with DMD.

Restrictions on non-Clifford fault tolerance and ruling out beyond-SQL quantum metrology

Constantin J. Cedillo-Vayson de Pradenne

Mentors: Jordan Cotler and Robert Huang

Quantum metrology promises a quadratic speedup over the standard quantum limit (SQL), but signal-aligned noise is expected to preclude this advantage in realistic settings. A potential route around known no-go results is to encode the sensors in a quantum code where the physical signal acts transversally as a logical gate. Understanding restrictions on transversal non-Clifford gates is therefore central to both quantum metrology and fault-tolerant quantum computation. Here, we prove such restrictions and apply them to transversal sensing. For any stabilizer code of distance $d \geq 3$ supporting a transversal logical action in level D of the Clifford hierarchy, every stabilizer generating set must contain a check of weight at least 2^D . Moreover, any r -level concatenated realization satisfies $r \leq \lfloor \log_2 n/D \rfloor$, forcing $r=1$ and ruling out concatenation when applied to beyond-SQL metrology. We then show that transversal single-qubit rotations by a small angle θ can only induce a nontrivial logical action on an n -qubit code if its checks include irreducible stabilizers of weight $\Omega(1/(n|\theta|^2))$. Here, many single-qubit errors commute with every stabilizer or logical Pauli below this weight and are only detected by a high-weight check, so their syndromes cannot be fault-tolerantly reconstructed from low-weight normalizer measurements. Since beyond-SQL transversal sensing requires $|\theta| = o(n^{-1/2})$, the weight of checks required for syndrome extraction diverges with n . Finally, we prove a broader metrological no-go theorem that avoids the assumptions of the quantum Cramér-Rao bound: constant-strength signal-aligned noise rules out any asymptotic advantage over the SQL in AC or DC sensing, even with biased estimators, nonstabilizer or approximate encodings, quantum memory, intermediate measurements, or adaptive control.

Confirming functionality of a single pixel bolometer for use in shearing interferometry

Ian J. Chang

Mentor: Sunil Golwala

Over this presentation, I tested the functionality of a bolometer that would be used to measure the uniformity of the surface of the Leighton Chajnantor Telescope using a tactic called shearing interferometry. I first go over the background of the telescope and shearing interferometry, which requires a single pixel bolometer and mirrors. I then describe the structure of the bolometer and its surrounding dewar, which was built to use liquid helium to keep the bolometer at a temperature under 1 kelvin. I go over the tests that I ran on testing the functionality of the dewar and its ability to cool the bolometer for an extended period of time. I first start with testing the dewar's ability to hold vacuum, necessary to thermally isolate the bolometer and the cold liquid nitrogen and helium inside the dewar. I then test the behavior of the dewar under liquid nitrogen, and I finally go over the response of the dewar under operating conditions, with liquid nitrogen and helium in separate tanks to create a cryogenic environment to sufficiently cool the bolometer. I go over the results of the tests and look ahead towards how the bolometer will be used to help the construction of the Leighton Chajnantor Telescope in Chile.

Cell-type-specific germline immune variants as predictors of immunotherapy response

Anusha K. Chatha

Mentors: Hannah Carter, Shinsuke Shimojo, and Noa Popko

Inherited genetic variation in the immune system influences cancer progression and patient response to immunotherapy, but the specific immune cell populations mediating these effects remain unclear. This project aims to investigate how germline immune regulatory variants contribute to immunotherapy response by integrating cell-type-specific expression quantitative trait loci (eQTLs) and chromatin accessibility data to identify the immune cell populations most relevant to treatment response. eQTLs from the Database of Immune Cell Expression (DICE) and OneK1K datasets were filtered to retain variants associated with immune regulatory genes, and linkage disequilibrium expansion was performed to identify linked regulatory variants. The expanded variants were integrated with immune cell-type-specific chromatin accessibility data generated by the assay for transposase-accessible chromatin using sequencing (ATAC-seq). Statistical testing of accessibility

across stimulated immune cell types identified 647 variants located in regions with cell-type-specific patterns of chromatin accessibility. After matching these variants to an immunotherapy-treated patient cohort and filtering, 614 variants were retained for machine learning analysis. XGBoost models were trained using all variants and cell-type-specific variant sets to predict immunotherapy response. Cross-validation showed modest predictive signal from germline variation alone, with a tuned all-variant area under the receiver operating characteristic curve (AUC) of 0.611 and cell-type-specific models with the highest cross-validation AUCs reaching approximately 0.60. Model performance was positively associated with the number of variants included in each cell-type feature set, suggesting that differences in feature-set size may contribute to observed differences in performance between cell types. These results suggest that cell-type-specific germline regulatory variation may be relevant to immunotherapy response, but further validation is needed to determine whether particular immune cell populations provide unique predictive information.

On-policy self-distillation for physically plausible video generation

Hongqiao Chen

Mentors: Georgia Gkioxari and Ziqi Ma

Video generative models have shown potential as world models, predicting futures from images or text. Yet they routinely generate physically implausible videos, even though almost every video they were trained on is real and physically valid. We hypothesize that this is because flow matching is mode-covering, so a well-fit model spreads probability into the thin regions between modes, and in video what lies between two plausible rollouts is not a third plausible rollout but an incoherent one. In this work, we propose a novel on-policy self-distillation method that sharpens the model toward its own high-density regions and show promising initial results.

Neural operators and stochastic methods for PDE-based navigation

Michael H. Chen

Mentors: Anima Anandkumar, Chuwei Wang, Hong Chul Nam, and Xi Deng

PDE-based navigation is a technique that produces global fields which represent environment geometry. The gradients of these global fields define local motion towards a target across complex domains, such as robot control with many degrees of freedom. However, due to the curse of dimensionality, conventional numerical methods turn prohibitively expensive to solve higher-dimensional configuration spaces. Accordingly, neural operators provide an elegant, scalable way to learn mappings from boundary conditions and domain geometry to these fields. In particular, we combine stochastic representations of elliptic PDE solutions with neural operators to learn continuous navigation fields. Deterministic numerical solver fields are used as references, whilst GPU-accelerated stochastic computation generates training data from maps diverse in geometry and dimension. Experiments show that local structure, boundary treatment, learned correction, and local control improve field quality and navigation success across varied environments. By amortizing stochastic PDE estimation across a family of environments, learned operators produce navigation fields substantially faster than conventional full-domain numerical methods.

Initial data dependence of direct waves in binary black hole mergers

Patrick P. Chen

Mentors: Yanbei Chen and Andrew Laeuger

Black hole mergers consist of an inspiral, merger, and ringdown phase. The ringdown consists primarily of quasinormal modes (QNMs), decaying sinusoids resulting from free oscillation of the remnant black hole. The merger consists of the direct wave, the last part of the wave sourced by the plunging particle and can carry signatures of the black hole horizon. QNMs have larger amplitudes than the direct wave, so we apply a frequency domain rational filter, which cancels out the QNM poles, and a full filter based on the response term in the drive-response decomposition of a perturbation waveform, to numerical relativity simulations from the Simulating eXtreme Spacetimes (SXS) catalog

in order to extract the direct wave. We find a remnant wave whose complex frequency evolves with increasing decay rate and whose properties are primarily dependent on the spin of the remnant black hole. The direct wave can potentially be used as an additional method for inferring remnant black hole properties used alongside standard QNM analysis.

Solving the Navier-Stokes equation for electron fluids of finite geometries

Runqin Chen

Mentors: Marco Bernardi and Tommaso Chiarotti

Electron hydrodynamics is a transport regime in which momentum-conserving electron-electron scattering dominates over momentum-relaxing scattering, allowing charge carriers to behave collectively as a viscous fluid. For a 2D material, it is required that $lee \ll W$ and $lee \ll l_{MR}$, where lee is the length scale of electron-electron interaction that conserves momentum, W the sample width, and l_{MR} the mean free path of momentum-relaxing scattering. This project aimed at developing a computational framework for simulating the regime in finite 2D materials. The model was built to first reproduce the results in graphene/hexagonal boron nitride (hBN) observed by Sulpizio *et al.* A steady-state Navier-Stokes equation containing Lorentz force was solved in OpenFOAM using parameters and boundary conditions inferred from the literature. The transverse Hall response was obtained from the simulations, and the normalized curvature of Hall field was compared with the experimental values. The simulated values followed the published trend in the literature, with deviations ranging from 3.5% to 7.4% across the carrier densities considered. The result provided a benchmark for the approach, supporting the plan to extend to anisotropic materials.

Development of a sample mount and high dynamic range power meter for ytterbium-based quantum interfaces

Sophia X. Chen

Mentors: Andrei Faraon and Erin Liu

Quantum networks rely on interfaces between optical photons and matter-based qubits. Rare-earth-ion solid-state systems, particularly ytterbium in yttrium orthovanadate ($^{171}\text{Yb}^{3+}:\text{YVO}_4$), offer a promising platform due to their long spin coherence lifetimes and fast microwave control. This project develops a cryogenic compatible sample mount for the $^{171}\text{Yb}^{3+}:\text{YVO}_4$ crystal, and a cost-effective, high dynamic range laser power meter that is optimized for the zero-phonon optical transition between the $^2F_{7/2}(0)$ ground state and the $^2F_{5/2}(0)$ excited state of the ytterbium ion at 984.5nm. This project aims to improve the operational efficiency of the experimental setup, by improving the accessibility of the $^{171}\text{Yb}^{3+}:\text{YVO}_4$ crystal and developing a highly compact and reproducible power meter. These developments provide a practical foundation for advancing ytterbium-based quantum network research, enabling more efficient characterization and control of the $^{171}\text{Yb}^{3+}:\text{YVO}_4$ spin-photon interface for future quantum networking applications.

Discovering interpretable visual features in high-fidelity generative models

Trevor B. Chen

Mentors: Pietro Perona and Ilona Demler

Understanding and controlling the visual concepts represented inside generative models requires features that are both interpretable and causally testable. This project studies whether sparse feature methods can decompose a compact image representation into human-interpretable and editable factors. I use a pretrained FFHQ-128 diffusion autoencoder as an experimental substrate because its semantic representation is a single 512-dimensional vector with no spatial or token-position dependence, unlike the spatially distributed residual streams of a DiT, while its decoder still produces realistic faces. I verified that the representation is deterministic, expanded the code corpus from 70,001 FFHQ examples to 232,771 FFHQ and CelebA examples, and trained TopK sparse autoencoders with different dictionary widths and active-feature budgets. Candidate features are evaluated by modifying their coefficients and passing the edited representation through the actual image generator, rather than by inspecting autoencoder reconstructions alone. This distinguishes features that correlate with an image property from features that causally control it. Initial scalar SAE experiments show that increasing data and capacity improves reconstruction, but feature usage and foreground/background

entanglement remain important concerns. To test whether visual concepts occupy multidimensional subspaces, I added a Block-Sparse Featurizer comparison that selects coordinated blocks of latent directions. The goal is to discover stable edits to glasses, hair, face shape, lighting, and background, and to determine whether structured sparse representations provide cleaner image features than scalar SAEs.

Systematic study of the long-term behavior of hot electrons surrounding the supermassive black hole

Wenda Chen

Mentors: Fiona A. Harrison and Xiurui Zhao

Most of the X-ray photons in the extragalactic universe are emitted by accreting supermassive black holes (SMBHs), known as active galactic nuclei (AGN). The intrinsic X-ray emission from AGN is primarily produced in a region surrounding the SMBH, known as the corona. Understanding the physical properties of this cosmic X-ray emitter is therefore essential for revealing how SMBHs release their accretion energy. However, the physical properties of AGN coronae remain poorly understood. Current coronal models predict that the coronal temperature should evolve with luminosity. In this work, we perform the first systematic study of the evolution of coronal temperature with luminosity. We carefully select a sample of 10 AGN whose coronal temperatures can be well constrained with current X-ray telescopes and that have multiple epochs of high-quality X-ray observations with the NuSTAR space telescope. This project enables us to systematically probe the evolution of coronal temperature for the first time. Our results provide important new insights into the physical mechanisms governing AGN coronae and have significant implications for current theoretical models of the corona.

Modeling experimental asset bubbles through Large Language Models with persona optimization

Yi-Jhe Chen

Mentors: Colin F. Camerer and Cady H. Ngo

Experimental asset markets are an effective way to study human trading behavior. With recent advancements in Large Language Models (LLMs), simulating human experiments with these models has become a new paradigm in social science. This project aims to build a robust environment with LLM agents in the experimental asset market, providing a proxy for studying human behavior in stock markets. By parametrizing and clustering experimental human data, we identified four types of traders in the experimental market. An optimizer LLM was then adopted to optimize natural-language personas, prompting LLMs to behave in a way that resembles humans. Our empirical results show that LLM agents are capable of reproducing human behavior and market paths with our optimized personas, demonstrating the promise of explaining human behavior with LLM agents.

Magnetic order in the bilayer nickelate $\text{La}_3\text{Ni}_2\text{O}_7$ from a two-orbital model

Yufan Cheng

Mentors: Garnet K. Chan, Petra Shih, and Yuhang Ai

The bilayer nickelate $\text{La}_3\text{Ni}_2\text{O}_7$ becomes a high-temperature superconductor under pressure. In many such superconductors, superconductivity is closely tied to magnetism, but the magnetic order of $\text{La}_3\text{Ni}_2\text{O}_7$ under high pressure has not yet been determined experimentally. At ambient pressure, several experiments point to a “double spin stripe” order, whose pattern repeats every four Ni atoms. Most previous calculations for the high-pressure phase, however, either did not compute an ordered ground state or used simulation cells too small to contain this pattern. Here we study a two-orbital model of the high-pressure phase with density matrix embedding theory, a method that treats the interactions between electrons more accurately within a simulation cell large enough to hold the double spin stripe. We find that the double spin stripe is the magnetic ground state, with antiparallel spins on the two layers, within relevant parameter range. Moreover, the two Ni orbitals behave very differently: one becomes nearly half filled and its electrons tend to localize, while the electrons in the

other remain mobile, a behavior known as orbital-selective Mott behavior. The double spin stripe is selected by electron correlation: it does not appear in a simpler mean-field treatment. Also, several magnetic orders lie close in energy, suggesting that the magnetic interactions within each layer are frustrated.

Physics is the best teacher: Consistency learning for time-invariant operators of chaotic dynamics

Lufang A. Chiang

Mentors: Anima Anandkumar and Jiachen Yao

Accelerating the prediction of long-term behavior in chaotic systems is crucial in scientific computing. However, existing methods rely on numerical solvers or autoregressive models that advance one small step at a time, which makes long horizons expensive. We instead view this problem as learning the system's time-invariant evolution operator, which jumps the state across a large time span in a single evaluation. To this end, we derive the consistency equations a time-invariant operator must satisfy, with differential and compositional objectives in physical time. These equations also tie the operator to the instantaneous dynamics, so we extend them to develop physics-distilled consistency learning. Across five chaotic systems, we find that physics-distilled consistency makes both short-term trajectories and long-term statistics more accurate. The learned operator survives temporal extrapolation and spends ten times fewer evaluations than autoregressive rollout, offering an efficient route to the long-term simulation of chaotic dynamics.

Helping Euclid see dark matter halos of galaxies

Emily Chien

Mentor: Andreas L. Faisst

The Euclid space telescope will find over one hundred thousand strong gravitational lens systems (consisting of a lens and a background source). However, most lens systems are blended at the resolution of Euclid's near-infrared slitless spectroscopy. This prevents us from deriving the deblended spectra and spectroscopic redshifts of lens and source needed to model each lens galaxy's total mass, including dark matter. This project develops the first module of the Strong Lensing Extraction Tool for Euclid (SLETE): a machine-learning framework that predicts a galaxy's most likely rest-frame spectrum given only its Euclid colors. I trained a self-organizing map (SOM) on $\sim 10,000$ SDSS galaxies with clean photometric and spectral data ($z < 1$), supplemented by GALEX and 2MASS data for ultraviolet and near-infrared coverage. By tracing outliers, I diagnosed and reduced systematic errors in the model — for example, flagging galaxies susceptible to the color-redshift degeneracy problem of where its spectral features fall between filter bands, and mapping quiescent and star-forming galaxies separately to further improve accuracy. The resulting pipeline matches any galaxy's Euclid colors to the SOM cell carrying its most-likely redshift and spectrum.

Predicting isotope effects with orbital-based deep learning

Seokhyun Chin

Mentors: Anima Anandkumar, Beom Seok Kang, John M. Eiler, and Simon Andren

Isotope effects play an important role in providing a mechanistic understanding of chemical reactions, but the exact quantification of isotope effects is limited by the high computational cost of quantum chemistry methods. Previous work has demonstrated notable success in predicting equilibrium isotope effects using AI, but it is only able to predict values for H, C, N, and O and excludes sulfur-bearing compounds. In this work, we demonstrate the capabilities of AI to predict equilibrium isotope effects of sulfur-bearing compounds and kinetic isotope effects. A chemical understanding of Sulfur isotope effects is first built by comparing the chemical environments of each sulfur atom and corresponding isotope effect values, and an SE(3)-equivariant graph neural network is trained to predict isotope effects of sulfur-bearing compounds and kinetic isotope effects. The model demonstrates high

performance in both cases, predicting sulfur isotope effects within experimental accuracy of 1‰ and kinetic isotope effects of heavy atoms within 5 ‰. The trained models can serve as a valuable tool to be applied to reaction mechanism elucidation of small organic molecules at a much lower computational cost compared to conventional quantum chemistry methods.

Evaluating the effectiveness of different transposon systems for genomic integration in chicken embryos

Pranati N. Chintada

Mentors: Marianne Bronner and Junpeng Gao

Developmental biologists seek to understand the molecular factors regulating embryonic development. Although CRISPR/Cas9-mediated loss-of-function experiments can determine the roles of molecular factors, they do not provide stable, long-term labeling of perturbed cells. To address these limitations, this project evaluates three transposon systems (AcDs, PiggyBac, Tol2) for stable genomic integration to enable long-term lineage tracing following loss-of-function experiments. These systems mediate genomic integration through a transposase enzyme that inserts DNA flanked by transposon sequences into the genome. Chicken embryos were electroporated with each transposon system, and fluorescence was monitored over one week. The AcDs system sustained weak fluorescence for 3 days. The PiggyBac system was not evaluated due to unsuccessful cloning of its transposon. The Tol2 system maintained strong fluorescence over 6.5 days. In conclusion, Tol2 produced the most sustained fluorescence, supporting its use for future lineage-tracing studies. Future research will focus on using Tol2 in loss-of-function experiments to study molecular factors involved in the development of the peripheral nervous system.

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.

Effects of dimensionality and forcing conditions on vortex formation and interactions

Minseung Choi

Mentors: Tim E. Colonius and Caroline Cardinale

Vortex dipoles and vortex rings are coherent flow structures that form in two-dimensional and three-dimensional flows and can interact through behaviors such as leapfrogging and merging. Although vortex dipoles are often treated as two-dimensional analogues of vortex rings, the effects of dimensionality on their formation and interactions are not fully understood. This project uses numerical simulations to investigate how forcing conditions influence isolated vortices and periodically

generated vortex trains. Localized body forcing was used to generate vortex dipoles, and numerical methods were developed to identify the leading vortices, separate them from trailing vorticity, and track their motion. Short forcing durations produced compact dipoles whose motion agreed closely with an idealized point vortex model prediction. Longer forcing durations produced increasingly prominent trailing vorticity and reduced the rate at which circulation accumulated in the leading vortex. Higher Reynolds numbers reduced the loss of circulation caused by viscosity. Dipole trains exhibited leapfrogging at all forcing frequencies examined, indicating that more widely spaced vortices should be studied to identify transitions between interaction regimes. Future work will extend this research to three-dimensional vortex rings and ring trains to determine which behaviors are preserved and which differ across dimensionality.

Mechanical and rheological characterization of guayule latex systems

Zoe R. Chou

Mentors: Julia A. Kornfield, Florence Müller, and Siddharth Premnath

Guayule (*Parthenium argentatum*) is a drought-tolerant shrub that produces hypoallergenic natural rubber (NR) and may act as a domestic hedge against global NR supply disruptions. During traditional guayule latex extraction, resin remains in the resulting latex and may affect its properties and the performance of films produced from it. The Kornfield lab developed a novel approach to extract resin-then-latex sequentially, and this project aims to investigate how the removal of resin affects the flow behavior and mechanical properties of guayule latex. Conventionally extracted guayule latex, which still contains resin, was used to develop procedures for characterizing the material prior to testing latex that went through the resin-then-latex extraction process. The flow behavior of the latex was measured over a range of strain rates and concentrations using a rheometer, while the mechanical properties of latex films were evaluated through tensile testing using an Instron machine. These experiments provide a framework to evaluate the quality of latex produced through sequential extraction processes, towards the larger goal of enhancing the overall value of guayule as a commercially viable crop.

Fiber entanglements effects mechanical behavior in fiber bundles: Experimental observations and analysis

Malia A. Christy

Mentors: Angkur Shaikeea and Rosan Naik

Entanglement in fibrous materials strongly influences their mechanical performance, yet quantitatively characterizing this relationship remains challenging. Understanding how fiber-scale structure translates to macroscopic mechanical behavior requires a multidisciplinary approach linking three-dimensional microstructure to mechanical performance. This study investigates how fiber entanglement affects the mechanical strength of fiber bundles using nylon monofilament as a model fibrous system. Controlled entanglement patterns are introduced into fiber bundles at low, medium, and high levels of entanglement and characterized using X-ray computed tomography (XCT). Three-dimensional fiber trajectories are reconstructed from XCT data using spline-based fits and analyzed computationally to quantify geometric and topological measures of entanglement, including slack length and fiber interactions. These reconstructed structures provide a framework for relating fiber-scale architecture to macroscopic mechanical behavior. Tensile testing is then used to characterize the force-displacement response and failure strength of bundles with varying degrees of entanglement. By combining three-dimensional imaging, computational trajectory reconstruction, quantitative entanglement metrics, and mechanical testing, this work establishes a methodology for investigating how fiber entanglement influences the mechanical performance of fibrous materials.

The dynamical consequences of stellar spin-orbit misalignment in exoplanet systems

Daria Alice Ciobanu

Mentors: Konstantin Batygin and Luke Handley

The formation and migration histories of close-in planetary systems remain longstanding problems. The relative orientations of the angular momentum vectors associated with stellar rotation and planetary orbits have emerged as a key diagnostic in differentiating competing theories of their

arrival. Among warm sub-Saturn systems, stars with several transiting planets tend to be aligned with their planets, while systems with only one known transiting planet more often show misalignment. This pattern is frequently interpreted as evidence that single-transiting systems underwent a dynamically violent formation or migration history. In this work, we investigate an alternative explanation: a primordial misalignment between the stellar spin axis and an initially flat planetary system may itself excite mutual inclinations and cause a multi-planet system to appear single-transiting. We simulate the primordial evolution of TOI-5398, a well-constrained representative system, and examine how the dispersal of its disk affects the surviving inclination. Preliminary results show that a tilted, rapidly rotating young star can excite enough mutual inclination to reduce the likelihood that both planets transit. If this mechanism operates broadly, the association between single-transiting systems and stellar misalignment is therefore an observational correlation arising from the misalignment itself, as opposed to a tracer of a distinct migration history.

Neural ODE across the cylinder wake bifurcation

Nicholas A. Ciordas

Mentors: Steven Brunton, Soon-Jo Chung, and Mario Sinani

This work asks whether one data-driven reduced model can represent the two-dimensional cylinder wake on both sides of its primary Hopf bifurcation. The dataset contains 34 trajectories spanning 17 Reynolds numbers from 30 to 160. Across this sweep, the same learned vector field must reproduce decay toward the steady wake below onset and growth toward a stable limit cycle above it. Seventeen independently trained neural ordinary differential equations reproduce trajectories but give uneven estimates of the fixed-point growth rate. A single parameter-conditioned network instead learns the Reynolds-number dependence from the pooled dataset. In its best holdout configuration, the shared model reduces mean growth-rate error from 26.7% to 5.2% over $48 \leq Re \leq 90$, reaches 2.2% error at held-out $Re = 51$ and 5.2% at held-out $Re = 73$, and locates the measured critical Reynolds number, $Re_{c_1} = 44.2$, within 1.9%. These results show that parameter sharing can replace a collection of local models with one compact model that interpolates across the bifurcation.

Flow-DFL: Decision-focused learning for decision-making under uncertainty via energy-based conditional flow matching

Tommaso Colombo

Mentors: Steven H. Low and Han Xu

Decision-making under uncertainty often optimizes for improving prediction accuracy to improve decision quality rather than to making high-quality downstream decisions. This project studies that “predict-then-optimize” gap in the setting of virtual electricity trading. The project compares a classical DPDS non-neural virtual bidding benchmark against four neural decision-focused policies. These policies each use the spread predictions generated by the prediction models to produce optimal trading decisions using either an MLP or an energy-based conditional flow matching (ECFM) decision model. Policies are then evaluated using validation total cumulative profit, Sharpe ratio, spread MSE and latency. Across 10 random seeds, the MLP + ECFM pipeline achieved the highest cumulative validation profit out of all the models, while FM-TS Direct achieved the highest average Sharpe ratio. The results suggest that flow-based models generally make better decisions when the prediction quality is the same, but also that the benefits are sensitive to model architecture, the initial seeds and the different regularization choices. Furthermore, the spread and Sharpe regularized model performed better than the non-regularized model in a single seed run. This suggests that optimizing for downstream decision making rather than prediction accuracy could be more important in making more optimal decisions under uncertainty.

Molecular dynamics simulations for nanoconfined ion transport

Carl J. Crum

Mentors: Kara D. Fong and Edward Chen

Upscaling global lithium harvesting for clean energy demands advanced separation technologies like nanofiltration, which rely on selective ion transport through nanometer-wide channel pores. Inside nanoscale confinement, intense ion-ion interactions and anisotropic dielectric screening invalidate

classical dilute-solution models, compelling the use of an Onsager transport framework to account for correlated ion motion. In this project, we performed molecular dynamics simulations in LAMMPS to compare unconfined bulk liquids against nanoconfined electrolytes across single-salt and mixed-salt solutions containing monovalent and multivalent cations. By calculating Onsager transport coefficients from particle displacements obtained through our MD simulations, we evaluated fundamental transport metrics, such as ionic conductivity, transference numbers, self-diffusion coefficients, and solution ionicity, while benchmarking bulk simulation densities against experimental literature. These results establish a control dataset and refine non-equilibrium transport theories, providing essential physical insights to optimize membrane selectivity for sustainable lithium extraction.

Resolution agnostic 3D MRI reconstruction using neural operators and diffusion priors

Ram Daftari

Mentors: Anima Anandkumar, Jiayun Wang, and Valentin Duruisseaux

Deep-learning-based 3D MRI reconstruction is fundamentally limited by the GPU memory and compute costs required for full-volume training. While training at low resolutions and running inference at higher ones offers a natural workaround, standard neural networks are resolution-specific and performance drops sharply outside their training settings. To address this, we propose a coarse-to-fine, resolution-agnostic framework for accelerated 3D MRI. Our method can be trained on memory-feasible, low-resolution volumes and applied zero-shot to higher-resolution data.

At its core, the framework utilizes neural operators that allow sampling at arbitrary resolutions. Across standard 3D knee and brain MRI datasets, our approach significantly outperforms classical, CNN, and transformer-based baselines, demonstrating substantial improvements, particularly in zero-shot, higher resolution transfer scenarios.

Finally, we extend this resolution-flexible approach to generative diffusion priors. Standard fixed-resolution diffusion models typically used in posterior-sampling are replaced with discretization-agnostic models. With these two methods, we aim to chart a highly scalable, memory-efficient path toward resolution-agnostic 3D MRI.

Hunting for the most extreme nuclear transients

Grace Daja

Mentor: Matthew J. Graham

Extreme Nuclear Transients (ENTs) are a new subclass of the most energetic ($\geq 10^{52}$ ergs) and most luminous ($\lesssim -23$ mag) Ambiguous Nuclear Transients (ANTs). These ENTs exhibit a rapid rise time (~ 150 days) and power-law decay in their lightcurves, thought to be the product of a tidal disruption of a high-mass star with a supermassive black hole. To identify these rare events, we have conducted a preliminary phenomenological time series search and developed a 3-state Hidden Markov Model, which through custom transition masking and classification, outputs a labeled time series with Quiet, Firing, or Decaying states. Through the initial search, we present one new ENT, J025840.41-023134.5, with energetics of $\approx 10^{53}$ ergs. With this Model, we prepare for a large-scale search of Zwicky Transient Facility data to discover ENTs in the database and constrain the physical frequency.

Radio holography for the Deep Synoptic Array

Ngoc Thach Dang

Mentor: Stephen Padin

We developed an on-field holography method to measure surface accuracy of the Deep Synoptic Array (DSA) antennas using signals from bright astronomical objects, with a proof of concept established using DSA's 5m prototype antennas at the Owens Valley Radio Observatory. The DSA is a radio telescope array under construction in Spring Valley, Nevada consisting of 1650 6m diameter antennas with a broad frequency range of 0.7-2.0 GHz. Surface irregularities on a radio antenna distort its angular response, compromising its ability to perform high-fidelity observations, so accurate surface characterization is essential. Holography on small antennas at low frequencies poses particular challenges, primarily due to low signal-to-noise and severe radio-frequency interference from cell

phone communications. To address the SNR limitation, we use a long-duration holography technique in which integration time is increased for off-source grid points, with measurements split across multiple days.

Converting carbon dioxide into useful chemicals and fuels

Daniel C. Darahdgian

Mentors: Harry B. Gray and Jay R. Winkler

Lowering carbon dioxide emissions while also reducing demand for fossil fuels is made possible through the catalytic hydrogenation of carbon dioxide into hydrocarbons. However, few catalysts can hydrogenate CO₂ while not reducing the carbon-carbon double bonds characteristic of olefins or unsaturated hydrocarbons. Therefore, we studied a carbon nanotube supported iron-potassium catalyst, which has previously been reported to have high CO₂ conversion while maintaining great olefin selectivity. The catalyst was prepared by sequential loadings of iron and then potassium on multi-walled carbon nanotubes and catalytic tests were carried out on a fixed reactor bed setup. In our experiments, we found that catalyst loading had a significant impact on the product distribution. Lower catalyst loading resulted in primarily oxygenated products while higher catalyst loading resulted in mostly hydrocarbon products.

Discovery of thermal management fluids

Grace F. Davis

Mentors: Jeremiah Johnson, Geoffrey A. Blake, and Thejas Wesley

As AI computing demand grows, data centers must dissipate heat from high-performance chips, which are prone to temperature-accelerated failure. Conventional air cooling cannot manage the increasing thermal output of advancing chip technologies, while continued increases in chip power density motivate the exploration of next-generation thermal-management approaches. New cooling technologies are therefore crucial to the efficient and sustainable scaling of generative AI. Here, we investigate molecularly engineered immersion fluids for advanced thermal management. Through a combination of targeted small-molecule synthesis, NMR, and calorimetry, we investigate the structure-function relationships governing thermal-management performance. Together, these analyses establish a framework for evaluating advanced thermal-management fluids.

Designing an integrated, tuneable unitary in thin film lithium niobite for GBS

Tomas F. Davis

Mentors: Alireza Marandi and Ofir Yesharim

Gaussian Boson Sampling (GBS) is a prominent quantum sampling protocol for demonstrating quantum computational advantage, building upon the Boson Sampling framework proposed by Aaronson and Arkhipov in 2013. GBS processors have the potential to outperform classical supercomputers for specific sampling tasks and have also been proposed for applications including molecular docking, graph theory, and machine learning.

This summer project focuses on the design of a monolithic, integrated GBS processor based on thin-film lithium niobate (TFLN). While bulk lithium niobate has long been an industry standard due to its excellent electro-optic properties, TFLN provides strong optical confinement and low propagation loss, enabling the integration of squeezed-light sources, interferometers, and electro-optic components on a single chip. Using Lumerical simulations, we have designed an integrated unitary in TFLN as a component of a proof-of-principle, small-scale GBS demonstration. The design incorporates broadband, single-mode squeezed-vacuum light sources generated through type-0 spontaneous parametric down-conversion, curved waveguides and tapers demonstrating 2 dB/cm propagation loss for TM modes, and multi-mode interference couplers demonstrating 1.4 dB excess loss over an operational bandwidth of 200nm. Thermo-optic phase modulators will also be utilised, exploiting TFLN's large thermo-optic coefficient. The next objective will be to fabricate and experimentally characterise the design, providing a pathway towards a fully integrated, monolithic GBS processor.

Quantum representations for anomaly detection in active galactic nuclei light curves

Joaquin Gabriel Malachi A. de Castro

Mentor: Matthew J. Graham

Active galactic nuclei (AGN) are bright galactic centers powered by accretion onto supermassive black holes. While AGN naturally exhibit stochastic brightness variations, unusual deviations from this variability can indicate rare or physically meaningful changes in the accretion system. This project investigates when the inductive biases of different quantum representations become useful for AGN anomaly detection. In particular, we consider Quantum Variational Rewinding, adapt Quantum Reservoir Computing (QRC), and develop a novel Quantum Feature Map (QFM). We match each of these to a synthetic positive control deliberately sampled to show the structure each method is designed to represent. Our results demonstrate that these architecture-specific structures survived in real AGN from the Zwicky Transient Facility (ZTF) survey, and the top 5% of QFM anomaly candidates were 5.3 times enriched for the associated structure.

Non-deterministic categorical Lyapunov theory

Jacob de Juan Millon

Mentors: Aaron D. Ames and Joe Moeller

We extend an existing categorical framework for stability and Lyapunov theory of F-coalgebras to make it compatible with non-deterministic and stochastic systems. We relax some axioms of the current framework for additional flexibility. We amend the framework's limitations through two main approaches. The first approach redefines solutions flows through a lax commuting diagram (previously strictly commuting), equipping an object with a preposetal structure under which F-morphisms are preorder-preserving. A Lyapunov morphism for a system implies equistability: all solutions admit a common class K bound. The second approach uses the endofunctor's monad structure to build a functor between F-systems and id-systems, identifying a F-system f on a space E with an id-system $f^\#$ on the transition space FE . An F-algebra is used as an aggregator to transform the original norm for E and a Lyapunov morphism for f to corresponding ones for FE and $f^\#$. The resulting notion of stability for $f^\#$ depends on the aggregator chosen, providing flexibility in how stability is defined. We apply both frameworks to common classes of non-deterministic systems (Markov kernels and transition graphs), recovering known stability results (almost-sure, mean-square, L^p stability, etc.).

Reflection from the plunging region of accreting black holes: A systematic study of spin measurements

Mercan Demiroglu

Mentors: Fiona A. Harrison and Elias Kammoun

The spin angular momentum of black holes is a fundamental property which encodes information about their origins, evolution, and cosmological impact. X-ray reflection spectroscopy, i.e. modeling of the reprocessed high energy X-ray emission from the ionized accretion disks of accreting black holes, is a powerful tool for constraining spin, and over the past few decades it has become a commonly adopted method of estimating the spins of black holes in X-ray binaries. However, standard models typically assume that the reflecting medium, the spiraling accretion disk, is truncated at the innermost stable circular orbit (ISCO), thus neglecting contributions from the plunging region. We will present results from a new model that extends the emitting/reflecting region down to the event horizon considering the contribution of the plunging region. We investigate how the inclusion of this region affects the BH spin measurements. We show that neglecting the aforementioned plunging region contribution introduces systematic biases in parameter estimation, consistently leading to an overestimation of measured spin values, especially in systems with lower intrinsic spins. We investigate these biases in detail, identifying the regions of parameter space where the deviations are most significant and explore their physical origins. This has major consequences for the current best estimates of black hole spin distributions across the entire mass scale.

Investigating the nature of an unusual supernova: Spectroscopic classification, light curve modeling, and dust formation

Indie E. Desiderio-Sloane

Mentors: Mansi M. Kasliwal and Sam Rose

This project investigates the unusual supernova ZTF25aatclux, a transient that has received extensive follow-up observations through the RAT survey. The primary goal is to classify this supernova and determine its physical properties, including whether it is forming dust, the nature of its progenitor star, and the mechanism powering its unusual light curve. To accomplish this, spectroscopic observations and optical and infrared light curves were analyzed and compared with those of more typical supernovae. Light curve modeling was performed using tools such as REDBACK to constrain the physical processes responsible for the observed emission, while infrared observations were used to search for evidence of dust formation in the supernova ejecta. By characterizing ZTF25aatclux and comparing it with normal supernovae, this project aims to improve our understanding of supernova evolution, dust formation, and the diversity of stellar explosions, contributing to a broader understanding of the life cycles of massive stars and their impact on the interstellar medium.

Optimal follow-up of gravitational wave events with M4OPT and Rubin

Athena R. Devashish

Mentors: Leo Singer and Brad Cenko

The early detection of gravitational waves (GWs) and their electromagnetic counterparts has been crucial in improving constraints on astrophysical fundamentals. The Multi-Mission Multi-Messenger Observation Planning Toolkit (M4OPT) utilizes a Mixed Integer Linear Programming (MILP) scheduler to optimize kilonovae (KNe) detection strategies following GW events. The Vera C. Rubin Observatory (Rubin), equipped with the largest digital camera ever built, has a goal of exploring transients, which could allow for M4OPT to further optimize the probability of detecting GW counterparts with Rubin's technology. M4OPT was adapted for Rubin by implementing a ground-based observatory slew model and modifying the scheduler to follow the "gold" and "silver" strategies detailed in Rubin's 2024 ToO workshop report for NS-BH and BNS merger follow-up. Observing plans were generated for 75 sky map triggers from O5 using M4OPT's adaptive exposure time and Rubin's strategies. Rubin's gold strategy predicts 3 KNe with 4.75h allotted per trigger while M4OPT's adaptive equivalent predicts 3 KNe with an average of 0.65h per trigger. Rubin's silver strategy predicts 5 KNe with 1.45h allotted per trigger while M4OPT's adaptive equivalent predicts 6 KNe with an average of 0.86h per trigger.

Accurate greenhouse gas retrievals: Leveraging surface heterogeneity for correcting light-path variations

Leonardo Di Tommaso

Mentor: Christian Frankenberg

Accurate retrievals of greenhouse gas (GHG) emissions are essential to meaningfully informing policymakers on regional policies. However, the community has struggled to obtain retrievals of carbon dioxide (CO₂) with uncertainties below the real variation of gas concentration in the atmosphere (0.25% or 1-3 ppm). The main source of uncertainty in CO₂ retrievals comes from aerosol scattering. Different aerosol types carry different scattering behaviors coherently with Mie scattering theory, which makes it difficult to correct their impact on greenhouse gas retrievals.

To solve the problem, we explored the dependency of these biases on various aerosol parameters, surface reflectance and geometry by simulating satellite retrievals using the radiative transfer model vSmartMOM. We devised a pipeline that is able to estimate the impact of aerosols on the signal in simulated noise conditions similar to those of MethaneSAT, a high resolution GHG sensing satellite, by using the fact that atmospheric parameters can be considered roughly uniform over small areas, whereas the surface reflectance underneath is sharply varying.

These studies provide a model for current and future high resolution GHG remote sensing missions to successfully quantify and correct aerosol-scattering introduced bias, thus retrieving more accurate estimates of GHG concentrations.

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

Evan W. Dickson

Mentors: Glauco 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.

If a statistical model predicts that common events should occur only once in 10,000 elections, maybe it's the wrong model

Andrew Dou

Mentors: Jonathan N. Katz, Daniel Ebanks, and Gary King

Political scientists forecast elections, not primarily to satisfy public interest, but to validate statistical models used for estimating many quantities of scholarly interest. Although scholars have learned a great deal from these models, they can be embarrassingly overconfident: Events that should occur once in 10,000 elections occur almost every year, and even those that should occur once in a ten-quadrillion elections are sometimes observed. We develop the first generative accurate statistical model of all US federal legislative elections for the U.S. House (1946–2024) and Senate (1954–2024) along with the computational framework to train it. We find that even as the incumbency advantage rises and falls, voting patterns nationalize, and heresthetic maneuvering remake candidacies and campaigns, the fundamental quantities of interest related to American legislative elections remain constant over eighty years, including but not limited to the probability of incumbent defeat, the magnitude and dispersion of candidate effects, the midterm penalty, and the partisan bias of electoral maps. Most striking, we show that the same set of structural parameters describes both U.S. House and Senate elections, implying a universal data-generation process for American legislative elections in the postwar era.

Generalizing blue copper proteins as biomolecular spin qubits via room-temperature electron spin coherences in azurin

Demetrios A. Dresios

Mentors: Ryan G. Hadt, Arshiyah Alam Laaj, and Christian A. Totoiu


Molecular qubits can act as quantum sensors by exploiting the sensitivity of electron spin coherences to external perturbations. These systems can be rationally designed to display a desirable electronic structure and spin dynamics that enable interrogation of specific environments. Bio molecular qubits offer a promising platform for quantum sensing through native or recombinant expression of molecular quantum sensors that permit study of biological processes, especially those occurring at localized and interfacial regions in cells. Using time-resolved Faraday ellipticity/rotation (TRFE/R) spectroscopy, which enables all-optical measurement of electron spin decoherence at room temperature and in solution, this project establishes electron spin coherence in the type 1 blue copper protein azurin. This protein enables investigations into whether the reduction potential of this system correlates with its electron spin decoherence time. Such a relationship would provide a design principle for tuning biomolecular quantum sensors and insight into a potential link between spin coherence and electron transfer driving force.

Trimetallic nickel-based CO₂ reduction catalysts and progress towards a synthetic analogue to the [NiFe]-carbon monoxide dehydrogenase C-cluster

Vassilios P. Dresios

Mentors: Theodor Agapie and Leo A. Morag

Multimetallic catalytic centers are useful for converting inert molecules, such as CO₂, into value-added products. An enzyme expressed natively in some anaerobic microorganisms, carbon monoxide dehydrogenase ([NiFe]-CODH), contains a NiFe₄S₄ cluster active site at which CO₂-CO conversion at inoffensive temperatures and pressures. The catalytic mechanism of the cluster can be better understood *via* a synthetic analogue of this metal cluster, which would facilitate study of the electronic structure of molecules similar to the cluster at different points during catalysis. The binding of a 6,6'-di(N-imidazolyl)-2,2'-bipyridyl ligand **L1** to an iron-sulfur cubane cluster (Fe₄S₄R₄) may open a promising avenue towards assembling a synthetic model for [NiFe]-CODH. The possible nickel-iron cooperativity in this metallic active site moreover inspires other potential multimetallic catalysts for CO₂ reduction. Towards this goal, the proposed dianionic bis(carbene)bipyridyl ligand may also host trimetallic Ni₂M (M = transition metal) complexes for cleaving C=O bonds. The impact of the identity of M on the ability of this trimetallic complex to catalyze CO₂ reduction is investigated.



Cross-modality reasoning enabled by multimodal AI agents for diabetic progression inference

Chrystal Duan

Mentors: Jia Liu, Wei Gao, and Wenbo Wang

Artificial intelligence (AI) is increasingly being explored as a tool for analyzing and interpreting complex biomedical data. Diabetes studies produce heterogeneous patient data, from clinical diagnostics to islet hormone secretion dynamic profiles and to imaging of molecular features. However, integrating these distinct data modalities into a cohesive interpretation remains challenging. Large language model (LLM)-based AI agents may facilitate the integration of these diverse data sources and support the identification of cross-modal relationships relevant to diabetes research. Here, we developed a memory-augmented, multimodal multi-agent framework to systematically reason across heterogeneous patient data, identify relationships between clinical and functional phenotypes, and support researchers in interpreting complex patient cases. Extensive benchmarking was performed to characterize the performance and behavior of various agent architecture-LLM model

pairings across computational efficiency, diagnostic accuracy, reasoning quality, consistency, and diagnostic uncertainty. We further evaluated whether the multi-agent system could leverage previously characterized cases stored in a memory database to infer characteristics of a deliberately withheld modality from the remaining patient information with modality-masking experiments.

Characterizing dust from extragalactic Wolf-Rayet binary candidates with the James Webb Space Telescope

Mason G. DuBois

Mentors: Ryan Lau and Mansi M. Kasliwal

Wolf-Rayet stars are hot and aging stars that have lost their hydrogen layer, exposing the helium-rich interior. Many Wolf-Rayet stars exist in binary systems, and systems with a carbon-rich (WC) Wolf-Rayet star and a massive companion sometimes exhibit episodic dust production regulated by their orbit. While dust-producing WC binaries have been studied extensively in the Milky Way, few extragalactic candidates have been studied. We use the python package *space phot* to analyze JWST infrared imaging, performing photometry on two dust-producing WC candidates identified in SPIRITS, a time-domain survey with the Spitzer Space Telescope. We then test two approaches to model dust emission, DUSTY and modified black-body fitting. We obtained an optical spectrum of the candidate SPIRITS 14bqe with the Palomar 200-inch Telescope, which ruled out a WC designation. While not a WC system, SPIRITS 14bqe is shown to be dust-producing, and we are exploring alternative explanations of its nature. The second candidate, SPIRITS 16ln, is a confirmed WC star with an observed spectral energy distribution that can be fit by dust modeling techniques. Our analysis demonstrates that JWST observations are ideal for investigating dust production from extragalactic sources like SPIRITS 14bqe and SPIRITS 16ln.

Frustrated Lewis pair-mediated CO₂ activation for electrochemical carbon capture

Eesha C. Duvvuri

Mentors: William A. Goddard III and Chi Zhang

Electrochemical CO₂ reduction (ECR) offers a promising route for long-duration energy storage by coupling renewable electricity with the conversion of CO₂ into value-added chemical products. However, efficient CO₂ activation at low applied potentials remains a fundamental challenge. Herein, we investigate frustrated Lewis pair (FLP) adducts as metal-free platforms for low-potential CO₂ activation. Density functional theory (DFT) calculations were performed on a library of FLPs, comprising B/Al-centered Lewis acids each paired with P(tBu)₃ or PPh₃, in order to evaluate CO₂ activation through structural, thermodynamic, and electronic descriptors. All FLP adducts exhibit thermodynamically favorable CO₂ binding relative to the corresponding free acid-base pairs, with P(tBu)₃-containing systems displaying stronger CO₂ association. Upon one-electron reduction, O-protonation is favored thermodynamically over C-protonation, indicating a preferential pathway for subsequent CO₂-derived intermediate formation. Analysis of reduced radical anion structure shows that PPh₃-based systems exhibit greater radical conjugation and delocalization, whereas P(tBu)₃-based systems generally exhibit carbon-centered spin density. These results demonstrate that systematic variation of Lewis acid and base components provides a means of tuning CO₂ binding and the electronic structure of reduced FLP-CO₂ adducts. Current work focuses on establishing the mechanistic origin of the Al/B-dependent charge and spin distributions.

The influence of planetary mass on mean-motion resonance type in Kepler and TESS multi-planet systems

Maya D. Edelstein

Mentors: Fei Dai, Rob B. Phillips, and Zhecheng Hu

This work investigates how the type of mean-motion resonance (MMR) occupied by neighboring planets depends on planetary mass, using a combined sample of confirmed planets and planet candidates from the Kepler and TESS missions. Compact multi-planet systems frequently exhibit orbital period ratios near small-integer values, whose resonant structure can preserve dynamical signatures of planetary formation and evolution. We identify neighboring planet pairs across nearly 10,000 systems and estimate planetary masses from radii using a standard mass-radius relation. To

classify resonant pairs, we develop an adaptive tolerance scheme based on Farey-sequence mediant theory, in which the tolerance decreases as resonances become more closely spaced, avoiding the limitations of a single fixed tolerance. We find that lower-mass planet pairs preferentially occupy tighter resonances (e.g., 5:4 and 3:2), whereas higher-mass pairs preferentially occupy wider resonances (e.g., 2:1), with high statistical significance (Kolmogorov–Smirnov tests, $p < 0.01$). Standard resonance-capture theory, which depends primarily on the mass ratio of the planets, does not predict this dependence on absolute planetary mass. We therefore test a scaling relation predicted by resonant overstability and find that the observed trend agrees with the predicted direction, although its magnitude is weaker than predicted. An analytic two-body chaos criterion applied to systems with reliable eccentricity measurements further indicates that higher-mass planets at their observed 2:1 spacing are dynamically stable but would become unstable at tighter separations given their typical eccentricities. This provides a possible dynamical explanation for the observed mass-dependent resonance distribution. Our results suggest that the present-day distribution of resonance types retains signatures of mass-dependent dynamical processes, particularly resonant overstability, that shaped compact planetary systems during and after their formation.

Impact of phosphorylation on cardiac myosin heavy chain function in pythons and humans

Annabelle F. Eisner

Mentors: Leslie Leinwand, Michael H. Dickinson, and Dakota Hunt

Hypertrophic cardiomyopathy (HCM) is a disease that causes the heart's walls to thicken and decreases the volume of the heart chambers. It is typically characterized by hypercontractility, since the heart muscle must beat more forcefully than normal in order to pump an adequate amount of blood. In humans, the disease is irreversible. Pythons, however, experience analogous hypertrophy and increased contractility after feeding, but in a physiological and reversible manner. Previous work in the Leinwand lab has identified altered post-translational modification levels of myosin heavy chain in the python heart after feeding, associated with contractility changes. To study the remodeling of the python heart, we focused on three regions of python cardiac myosin that exhibit differential phosphorylation or acetylation after feeding. We also studied the analogous regions of human myosin, which often contain mutations in HCM-affected individuals. We used molecular cloning techniques to introduce single amino acid mutations at the regions of interest that mimic the negative charge and larger size of a phosphate group or the effect of acetylation. We also created control constructs that were phospho- or acetyl-null at these sites. Then we expressed these constructs by infecting cultured C2C12 cells with adenoviruses generated by the pAdEasy adenoviral vector system and purifying the resulting myosin via PDZ affinity chromatography. Finally, we used NADH-coupled ATPase assays on the purified myosin wild type and mutant constructs to study the impact of phosphorylation on the function of both human and python myosin. The S738D construct showed a significant increase in maximum ATP hydrolysis rate (V_{max}), consistent with the decrease in phosphorylation and slower ATP turnover that has been observed in post-fed python hearts. The S392D construct unexpectedly showed decreased ATP turnover rate, but did show a significant increase in binding affinity for actin that could explain this slower rate. Both mutations had significant effects on human myosin function, demonstrating the potential therapeutic power of python heart remodeling adaptations to halt diseases like HCM.

Beyond statistical isotropy in 3D galaxy surveys

Adam El-Hindi

Mentors: Olivier P. Doré and Henry Gebhardt

Statistical isotropy is the assumption that the Universe has no preferred direction and it underlies most analyses of large-scale structure. We develop an observer-centered, three-dimensional test of this assumption for wide-area galaxy surveys, with SPHEREx as the principal application. The galaxy density field is expanded in spherical Fourier-Bessel modes, and its full two-point function is decomposed into bipolar spherical-harmonic coefficients. The isotropic power spectrum appears as the monopole of this decomposition, while higher multipoles describe general direction-dependent correlations across angular and radial scales. We derive quadratic estimators for these coefficients and model their coupling through the survey geometry, radial selection, shot noise, catalog normalization, and redshift-space distortions. The estimator and covariance framework are validated on simulations

with an anisotropy signal injected. As a first survey application, we apply our estimator on the quadrupolar sector of the eBOSS DR16 quasar sample and find no significant departure from isotropy. This work provides the basis for a general multipolar analysis of SPHEREx, capable of testing cosmic isotropy in three dimensions and constraining a broad range of possible directional patterns in the galaxy distribution.

Trapping distinct states to characterize the MurJ-Lipid II complex

Kylin M. Elliott

Mentors: William M. Clemons and Grace F. Baron

MurJ is the essential lipid II flippase involved in peptidoglycan biosynthesis, the pathway responsible for cell wall formation that is conserved across bacteria. When bacteria cannot form a cell wall, cells will lyse from osmotic pressure, making MurJ an emerging antibacterial target. Due to conformational changes and subsequent substrate release, it has been impossible to capture the substrate-bound structure of MurJ or to fully understand its transport mechanism. To solve this problem, constructs with point mutation pairs were created, specifically engineered to form disulfide bonds that trap an inward-facing conformation of MurJ. We hypothesize that trapping this distinct state will allow for substrate binding in vitro and allow structural characterization of the complex. The primary methodology includes construct design and multi-step membrane protein purification, and evaluation of disulfide formation using assays with trypsin digestion and oxidative treatments. Disulfide bond formation in the constructs has been confirmed, and CryoEM analysis will be performed to evaluate the state of the proteins before moving forward with structural characterization of the lipid II-bound state.

High-dimensional communication via continuous wave lasers

Elizabeth C. Ellison

Mentors: Maria Spiropulu and Raju Valivarthi

High-dimensional quantum entanglement allows for denser information storage in quantum communication by utilizing a larger Hilbert space. Experimental implementations are severely impacted by channel loss, detector limitations, and system noise, which result in loss of information. In our work, we demonstrate long-distance distribution of high-dimensional time-bin entanglement with a continuous-wave laser that creates entanglement via spontaneous parametric down-conversion (SPDC) and unbalanced interferometers. The entangled photons were transmitted up to 200 km. This transmission was then improved by optimizing pair generation rate, dimension measured, and detector settings. The optimization showed ideal operational regimes that varied based on system loss. The results demonstrate long-distance high-dimensional entanglement distribution that is effective in realistic high-loss systems.

Entropic and enthalpic contributions to monovalent and divalent ion adsorption in electric double layers

Defne Nihal N. Ertugrul

Mentors: Kara D. Fong and Madeline Murphy

Positive ions adsorb at charged electrodes through enthalpic and entropic forces, but charge and size are hard to separate. Preliminary results in our group suggested sodium adsorbs for energetic reasons and calcium for entropic ones, though both differ in charge and size. To separate them, I built two artificial ions, one with sodium's size and double charge, one with calcium's size and single charge, giving four systems varying charge and size independently. I ran constant potential molecular dynamics in salt water between gold electrodes at 0 and 1 V across several temperatures, taking adsorption entropy from the free energy's temperature dependence. The adsorption entropy scaled with the first shell water shed, near 4 cal/mol/K per water, and charge set the amount carried and shed. Sodium was the least favorable entropically, losing entropy despite adsorbing strongly, consistent with the retained water being ordered, a shared penalty that sodium's small shedding fails

to offset. Because shedding and entropy both track charge, this attribution remains an interpretation. Confinement and interfacial ordering did not account for the trend. These results come from a single force field, calcium is undersampled at 298 K, and umbrella sampling confirmed sodium's free energy while calcium's stayed inconclusive.

Decoding the lanthanome: Medium-throughput characterization of protein affinity and selectivity towards lanthanides

Shrila Esturi

Mentors: Ariel Furst, William M. Clemons, Cooper Tezak, and Farid F. Khoury

The periplasmic protein CnrX in *Cupriavidus metallidurans* cells was engineered to have higher selectivity and binding affinity towards lanthanides for improved rare earth elements extraction. Specific residues in the binding pocket were mutated with MAX randomization to construct a saturation site mutagenesis library, and each variant was immobilized on a column with the SpyCatcher protein in 96-well plates. Through employing experimental techniques such as isothermal calorimetry and inductively coupled plasma mass spectrometry, the K_d constants and lanthanide selectivities of each variant were calculated. Using circular dichroism spectroscopy, nano differential scanning fluorimetry, and dynamic light scattering, the protein structure and dynamics were further studied. With this library, we aim to understand the energetics and thermodynamics of lanthanide binding interactions with proteins to develop a more sustainable approach to lanthanides mining.

Uniaxial strain tuning of Hall response in EuAuP

Mingjun Fan

Mentors: Linda Ye and Guoxin Zheng

EuAuP exhibits a giant Hall effect at zero magnetic field despite vanishingly small net magnetization, which has been proposed to originate from a real-space topological spin texture. This project investigates whether uniaxial strain can tune this response. Transverse and longitudinal voltages were measured from 2 to 10 K while sweeping a magnetic field along the crystallographic c axis and applying up to approximately $\pm 0.6\%$ nominal strain within the ab plane. The zero-field Hall voltage changed strongly up to 80 %, whereas the longitudinal voltage changed much less. The strain response depended strongly on temperature and reached its greatest magnitude near 5 K. Temperature-dependent measurements also showed that strain shifts a characteristic feature in the longitudinal voltage, suggesting that strain modifies a characteristic temperature associated with the magnetic state of EuAuP. Cooling the sample under compression further altered the low-temperature Hall response, although the reproducibility and microscopic origin of this history-dependent effect remain unresolved. These results demonstrate that small lattice distortions can provide substantial control over the Hall response of EuAuP, not only in the magnitude of the zero-field Hall voltage but also in the shape of its hysteresis loop.

Learning based resilient autonomy for robotic systems under actuation and perception faults

Baaqer M. Farhat

Mentors: Soon-Jo Chung and Mahdi Taheri

Robots operating in real environments must maintain stable behavior when perception becomes unreliable or task conditions change. This project studies methods that combine multimodal learning, reinforcement learning, and motion retargeting to improve robotic autonomy. The first part of the work uses a multimodal Transformer to process histories of state estimates, control inputs, and diagnostic signals. The learned representation updates model predictive control cost parameters and residual dynamics during operation. This approach adapts the controller instead of replacing it, which preserves the planning and safety structure of model predictive control. We test the method on a spacecraft simulator in the laboratory under controlled perception faults to evaluate its ability to maintain stable control when perception degrades.

The second part of the work applies robot learning to locomotion and whole body interaction tasks. Reinforcement learning policies are trained using about 4096 parallel simulation environments and produce joint position targets at 50 Hz. A locomotion policy is transferred from simulation to physical hardware and follows commanded forward, lateral, and turning velocities. Human motion data from a 52 joint model is also retargeted to a 31 degree of freedom robot and used as a reference for box manipulation and crawling policies.

Advancing the image simulator for the Ultraviolet Explorer (UVEX): Vignetting and red leak

Kenji P. Farrell

Mentor: Fiona A. Harrison

UVEX is an upcoming NASA space mission that will observe the universe at ultraviolet (UV) wavelengths. With the capability to detect sources 50 times fainter than previous missions, it will conduct the deepest UV survey of the sky. Equipped with Near-UV (NUV) and Far-UV (FUV) imaging units and a low-resolution spectrograph, UVEX will address a broad range of astrophysical questions, including the evolution of low-mass galaxies, early UV emission from cosmic explosions, and stellar astrophysics. UVEX is currently in its instrument design phase, during which the optics and detector configurations are being finalized, making realistic image simulators critical to this process. This summer, I have been working with the UVEX team to develop a new simulator for the NUV and FUV imagers using ScopeSim, an open-source instrument simulation framework. Building on a vignetting model developed during my previous SURF project, we refined the model by applying Gaussian smoothing and piecewise polynomial fits to convert noisy ray-tracing output into a continuous transmission map, which was then integrated into the simulator. We are now investigating red leak, in which out-of-band red photons pass through the UV filters and are detected as UV signal. By separating simulated stellar spectra into in-band and out-of-band components, we quantify how much of the recovered flux originates from genuine UV emission. The completed simulator will be integral to evaluating UVEX's expected performance and informing its final instrument design.

Surface chemistry behavior of carbon-based ablative thermal protective systems under hypersonic flow and temperature

Lucas J. Favretto

Mentors: Joanna M. Austin and William Feasey

Hypersonic vehicles experience extreme physical conditions when traveling at speeds of Mach 5 and above. Compressed air on the surface of the vehicle generates severe frictional heating and requires thermal protection systems (TPS) to protect the vehicle's structure from overheating. Carbon-based materials such as graphite are commonly used in ablative TPS because they can withstand high temperatures while undergoing surface reactions that dissipate heat. However, to improve the efficiency of such carbon-based models, engineers require a better understanding of the surface chemistry and reaction products generated when heated graphite interacts with high-enthalpy air.

The goal of this study is to develop a robust, resistively heated graphite model for investigating carbon surface chemistry and CN emission through spectroscopy in hypersonic test facilities. The current work focuses on improving previous heated-model designs to achieve graphite surface temperatures near 1500 K while limiting overheating, deformation, and structural failure in the surrounding components. Electrical and thermal models have been developed to evaluate candidate geometries, including a hollow half-cylinder and a blunted-wedge design, and preliminary bench-top tests using graphite strips are being used to examine the effects of geometry, electrical resistance, thermal mass, and contact conditions on heating performance. Because this study is still in progress, final experimental results have not yet been obtained. Future work will involve validating the final heated-graphite design before conducting hypersonic experiments and spectroscopic measurements of CN emission.

A manifold constrained Ensemble Kalman Filter algorithm for nonlinear dynamics

Ethan M. Feng

Mentors: Andrew M. Stuart, Bohan Chen, and Sangmin Park

The Ensemble Kalman Filter (EnKF) approximates the filtering distribution with a Gaussian ensemble updated via linear transport. While exact in the linear Gaussian mean-field limit, this approach breaks down for nonlinear dynamics. The ambient affine analysis map transports ensemble members away from the system's intrinsic attractor, causing systematic bias and an inflated posterior covariance.

We propose the Argus Ensemble Filter (AEF), a manifold-constrained variant that restricts the analysis update to evolve along the lower dimensional geometry of the forecast ensemble as opposed to the full ambient space. After a standard EnKF analysis, each member is projected onto a local polynomial chart learned from the forecast dynamics. We examine differing approaches to the local chart construction including a global chart and an adaptive moving least squares, additionally we introduce different projection variants that determine how the point cloud is assembled.

Evaluated on the Lorenz-63 system against a one million particle bootstrap particle filter, AEF significantly reduces the sliced energy distance to the true filtering distribution compared to EnKF. The manifold constraint corrects the systematic bias of the ambient analysis update by preserving inherent local curvature structure. These results suggest embedding EnKF updates within the intrinsic geometry of the forecast leads to a principled route to robust filtering in strongly nonlinear regimes.

Ultrashort pulses from an optoelectronically mode-locked laser on nanophotonic lithium niobate

Richard H. Feng

Mentors: Alireza Marandi, Rithvik Ramesh, and Yiting Wang

Mode-locked lasers (MLLs) are important sources in modern photonics because they produce periodic trains of ultrashort optical pulses. These pulses are useful for optical computing, time multiplexing, and driving nonlinear optical processes, while the corresponding frequencycomb enables applications in spectroscopy, precision timing, and ranging. The comb-line spacing, equal to the pulse repetition rate, provides a bridge between the optical and microwave frequency domains. Recent efforts have focused on developing integrated MLLs to enable compact, electrically driven photonic sources. Using a hybrid cavity formed by coupling a semiconductor gain chip to a thin-film lithium niobate (TFLN) waveguide, we demonstrate an electrically pumped integrated MLL with active intracavity loss modulation provided by a tunable mirror architecture. We generate pulses as short as ~ 300 fs with a coherent frequency comb spanning ~ 23 nm centered near 1550 nm, with over 1.5 W of on-chip peak power, over 400 mW of off-chip peak power, and over 1.5 mW of off-chip average power. We further demonstrate self-starting and self-sustained operation of the MLL as a coupled optoelectronic oscillator, with the laser generating its own RF drive signal at its repetition rate, eliminating the need for an external RF signal generator. In this configuration, we maintain ultrashort (~ 330 fs) pulses while simultaneously generating a 9 GHz microwave signal. Ongoing work focuses on fully stabilizing the frequency comb to ultrastable references through injection locking and other locking schemes, toward further development of the platform for low-noise microwave synthesis.

Transient ULX case study: M31 ULX-3 (LSXPS J004321.0+411833)

Ferdinand

Mentor: Murray Brightman

Ultraluminous X-ray sources (ULXs) are off-nuclear X-ray sources with luminosities exceeding the Eddington limit of typical stellar-mass compact objects. While most known ULXs are relatively persistent sources, transient ULXs emerge as a distinct population that offers unique insight into studying extreme super-Eddington accretion. We report the discovery and X-ray characterization of a newly identified transient ULX in M31. First detected serendipitously on 2021 June 9, the source reached a peak X-ray luminosity of $\sim 1 - 2 \times 10^{39} \text{ erg s}^{-1}$. Its spectral evolution suggests increasing advection at later epochs, opposite to the expected transition from a super-Eddington slim disk toward a standard thin disk during the decay of the outburst. The observed X-ray light curve can be

reproduced using the disk-instability model of Hameury & Lasota, with characteristic timescales of $\sim 60 - 85$ days. The inferred mass-accretion rate is super-Eddington for a $1.4 M_{\odot}$ neutron star but sub-Eddington for a $10 M_{\odot}$ stellar-mass black hole. Nearby optical sources from the *Hubble Source Catalog*, suggest a possible giant donor with an age of $\log_{10}(\text{age/yr}) \sim 8.5$ and a mass of $3.2 M_{\odot}$, when compared with stellar isochrones with an assumed $A_V = 1.58$. We estimate a transient ULX occurrence rate in M31 of $0.15 \pm 0.09 \text{ yr}^{-1}$. This rate is lower than those reported for other galaxies in previous studies. M31 ULX-3 expands the small population of known transient ULXs in M31 and highlights the importance of continued, uniform X-ray monitoring of nearby galaxies.

Resin development and monomer synthesis for additively manufactured, architected, glucose-responsive hydrogel microneedles

Lucas J. Flach

Mentors: Julia R. Greer and Analiese J. Wiedenbeck

The ultimate goal of this project is to design and fabricate architected microneedles for rapid transdermal delivery of glucagon and insulin for people living with Type 1 diabetes. Previous research has demonstrated mold-based fabrication of monolithic hydrogel microneedles, but these suffer from slow insulin release kinetics. This project aims to introduce micron-scale architected structures to enhance the surface area to volume ratio of the microneedles and accelerate insulin release. Our summer objective was to develop a resin formulation that (1) releases insulin in the presence of glucose and (2) can be additively manufactured into architected, mechanically stable cone structures. We accomplished this goal by optimizing the solubilizer composition of the resin and successfully fabricated structurally stable, architected microneedles. Additionally, we developed a scalable two-step synthesis of 2-aminoethyl acrylate hydrochloride (AAH), the resin monomer necessary for insulin binding. We synthesized 2.5 g of AAH, enabling high-throughput resin production for the fabrication of architected, glucose-responsive microneedles.

fMRI analysis of human first impressions in a four-dimensional trait space

Nicholas A. Flach

Mentors: Ralph Adolphs and Austin Brotman

Humans rapidly form first impressions when viewing unfamiliar faces, making judgments about traits such as warmth and competence. Since these social judgments influence many important real-world decisions, we seek to explore which brain regions are involved. We used functional magnetic resonance imaging (fMRI) to scan participants' brains while they viewed unfamiliar faces and periodically rated the similarity of pairs of faces. We sought to validate that their similarity judgments correspond with those from an independent group of rating participants. We then fit first-level encoding models on the fMRI data to identify brain regions whose activity is predicted by social traits and visual features, and fit decoding models to determine whether patterns of brain activity can accurately predict perceived social traits. Together, these analyses provide a greater understanding of how and where the human brain represents first impressions.

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.

Improving the sensitivity of NEMS resonators for mass and optical sensing

Ron A. Freeman

Mentor: Michael L. Roukes

Nanomechanical (NEMS) resonators achieve ultra-high sensitivity sensing by translating small shifts in resonator parameters like mass or temperature to shifts in resonant frequency. These systems have been developed for well over a decade and offer promising applications from massively parallel proteomics with single molecule resolution to a new class of uncooled, optomechanical IR detectors. However, they suffer from anomalous phase noise (APN) that limits ultimate detector sensitivity to orders of magnitude higher than the thermomechanical limit. Here, we present a holistic full system approach to characterizing, understanding and designing around APN. On the device side, we use theoretical modeling, and simulation to develop design intuition resulting in novel device design and a planned test campaign to understand and bound APN sources in silicon doubly clamped beam resonators. On the hardware side, we construct a dark port interferometer to drive readout noise to a minimum and implement real-time computational techniques to dynamically cancel noise processes.

Investigating thresholds of boulder mobilization in post-wildfire debris flows

Kirby Fung

Mentors: Michael P. Lamb and Emily Geyman

Post-wildfire debris flows are a significant hazard in mountainous regions, where intense rainfall on steep, burned landscapes can mobilize large volumes of sediment and boulders. Existing debris-flow prediction models estimate the likelihood and magnitude of debris flows but do not explicitly account for the conditions controlling boulder entrainment. Previous studies of boulder mobility have also largely focused on clear-water or simplified flow conditions, providing limited understanding of how sediment-rich debris flows affect boulder entrainment. This study investigates the physical conditions governing boulder mobilization following the 2025 Eaton Fire in the San Gabriel Mountains. High-resolution LiDAR (Light Detection and Ranging) surveys and optical imagery are used to identify and quantify individual boulder movement before and after storm events. Meta's Segment Anything Model 3 is used to identify boulders, and movement for each boulder is manually verified between surveys. For observed boulders, Shields stress is calculated as an indicator of the balance between flow forces and the resistance to boulder motion, and its relationship with channel slope is evaluated to identify factors influencing boulder mobility. These field-based results are compared with flume experiments conducted with clear water to quantify differences in entrainment between natural debris-flow and clear-water conditions.

Toward safe and efficient predictive control for spacecraft proximity operations

Deepak G

Mentors: Aaron D. Ames, Joe Moeller, Pedro Roque, and Yana Lishkova

Space missions involving close-proximity operations require efficient strategies that can simultaneously control position and attitude while optimizing the system's performance and respecting strict safety-critical constraints. We pose the problem of optimally controlling the coupled 6-DoF motion of the system with a tube-based convex Model Predictive Control (MPC) approach. For this purpose, we employ a state-of-the-art technique based on Difference of Convex (DC) functions to represent the coupled 6-DoF dynamics in a convex formulation. We derive an exact DC representation of the model and investigate learning-based methods based on Input Convex Neural Network (ICNN)

and SOS-convex polynomials. We exploit the structure of the system dynamics to reduce the size of the resulting optimization problem. We validate our approach within a layered architecture on a planar robotic testbed (the DISCOVER ATMOS free-flyer) that replicates spacecraft motion. Future work will focus on incorporating the model uncertainty and disturbances directly into the DC-MPC formulation.

SafePPO: Anytime constrained proximal policy optimization via safe gradient flow

Varun N. Gabbita

Mentors: Aaron D. Ames and Pol Mestres

Reinforcement learning (RL) policies for legged locomotion are typically trained without formal safety guarantees, allowing falls and instability to persist throughout training. This project presents SafePPO, an extension of Proximal Policy Optimization (PPO) using Reinforcement Learning-based Safe Gradient Flow (RL-SGF), an anytime-safe constrained optimization algorithm. At each training iteration, SafePPO solves a closed-form quadratically constrained update that balances task reward against a safety constraint, projecting the policy back toward safety whenever the constraint becomes active. We validate this approach first on a reduced-order linear inverted pendulum model, then on a full-order G1 humanoid walking task using center of pressure (CoP) as the safety constraint. Across both settings, SafePPO matches PPO's task performance while incurring significantly fewer safety violations, in simulation and on hardware. Additionally, we explore encoding reward terms as hard safety constraints, reducing the need for excessive reward tuning while preserving safety guarantees. Together, these results point toward a path for deploying learned locomotion policies with the rigor and predictability expected of classical safety-critical control.

Binodal: Towards an n-ary phase diagram generator for geochemical and geological applications

Vedant Ganesh

Mentor: Paul D. Asimow

Phase diagrams describe the behavior of chemical systems under conditions of thermodynamic equilibrium, summarizing stability across composition, temperature, and pressure space. Their construction involves finding the lower convex hull of the Gibbs free energy surface to determine the stable phase(s); this is numerically challenging when there are potentially many complex surfaces intersecting. Existing phase diagram construction software typically uses polynomial fits (e.g., Scientific Group Thermodata Europe polynomials) rather than physically motivated forms for the Gibbs free energy surface. We present Binodal, an open-source Python thermodynamic toolkit that calculates and visualizes binary phase diagrams using explicit thermodynamic models. Binodal uses thermodynamically consistent regular-solution-type models with pressure- and temperature-dependent Margules parameter in a user-transparent manner, enabling its use as a pedagogical tool as well as for lightweight thermodynamic calculations in petrology. We use this software to reproduce several common phase diagram topologies, such as eutectic, peritectic, and solvus diagrams, and investigate their topologies and thermodynamic potential functions in various spaces (e.g., Gibbs energy–temperature–pressure or internal energy–volume–entropy space).

FC-PINO: High-precision physics-informed neural operators via Fourier continuation

Adarsh Ganeshram

Mentors: Anima Anandkumar and Valentin Duruisseaux

The physics-informed neural operator (PINO) is a machine learning paradigm that has demonstrated promising results for learning solutions to partial differential equations (PDEs). It leverages the Fourier Neural Operator (FNO) to learn solution operators in function spaces and leverages physics losses during training to penalize deviations from known physics laws. Spectral differentiation provides an efficient way to compute derivatives for the physics losses, but it inherently assumes periodicity. When applied to non-periodic functions, this assumption of periodicity can lead to significant errors, including Gibbs phenomena near domain boundaries which degrade the accuracy of both function representations and derivative computations, especially for higher order derivatives. To overcome this limitation, we introduce the FC-PINO (Fourier-Continuation-based Physics-Informed Neural Operator) architecture which extends the accuracy and efficiency of PINO and spectral differentiation to non-

periodic and non-smooth PDEs. In FC-PINO, we propose integrating Fourier continuation into the PINO framework, and test two different continuation approaches: FC-Legendre and FC-Gram. By transforming non-periodic signals into periodic functions on extended domains in a well-conditioned manner, Fourier continuation enables fast and accurate derivative computations. This approach avoids the discretization sensitivity of finite differences and the memory overhead of automatic differentiation. We demonstrate that standard PINO fails (without padding) or struggles (even with padding) to solve non-periodic and non-smooth PDEs with high precision, across challenging benchmarks. In contrast, the proposed FC-PINO provides accurate, robust, and scalable solutions, substantially outperforming PINO alternatives, and demonstrating that Fourier continuation is critical for extending PINO to a wider range of PDE problems when high-precision solutions are needed.

Late-time HST photometry and X-ray analysis of a sample of Type II supernovae

Jiayi Gao

Mentors: Mansi M. Kasliwal and Wynn Jacobson-Galan

The shock produced by a Type II supernova collides with and interacts with the surrounding circumstellar medium (CSM). Probing the full spectral energy distribution provides valuable insight into the physical nature of this ongoing interaction and, ultimately, the properties of the progenitor star prior to explosion. We present optical (F555W) and ultraviolet (F275W) photometry obtained with the Hubble Space Telescope as part of Cycle 31 program 17506, combined with all available archival HST observations, including data obtained in additional filters. The raw observations are processed and drizzled using the *hst123* pipeline, after which point-spread function (PSF) photometry is performed with *DOLPHOT*. We construct late-time light curves for a sample of 30 Type II supernovae, enabling studies of their long-term evolution. Future work will include reducing X-ray observations for a subset of these supernovae and constructing the corresponding X-ray light curves. Together, tracing the luminosity evolution of these supernovae across optical, ultraviolet, and X-ray wavelengths will enable us to investigate shock-CSM interaction, constrain the CSM density, and infer the progenitors' mass-loss histories.

Heritable targeted DNA insertion using the R2Tg retroelement complex

Lucas N. Garcia

Mentors: Gözde S. Demirer and Kimberley T. Muchenje

Conventional techniques for inserting exogenous DNA into plant genomes using *Agrobacterium tumefaciens* result in random integration, often at multiple genomic loci. Targeted alternative platforms based on CRISPR/HDR have low editing efficiency, high off-target insertion events, or limited payload size. One emergent alternative technology is based on R2 non-LTR retrotransposons, which use target-primed reverse transcription to insert gene-sized DNA at the 28S/25S rDNA locus without generating double-strand breaks. Muchenje et al. have recently engineered the *Taeniopygia guttata* R2 protein (R2Tg) for targeted insertion at the 25S DNA site in plants, and demonstrated high insertion efficiency in *Nicotiana benthamiana*, *Arabidopsis thaliana* and *Solanum lycopersicum*. Here, to realize the platform's full potential, we go beyond transient systems and attempt to transform the R2Tg editor components into Micro-Tom tomato callus, detect insertions at the 25S rDNA target site via junction PCR and sequencing, and regenerate edited plants. Preliminary results suggest that continued RTg activity confers germline cytotoxicity, preventing retrieval of transformed seeds. To mediate this, we are testing R2g expression using the estradiol-inducible XVE system to narrow down the window and R2g expression level that allow retrieval of edited seeds. This work strives to establish stable, heritable R2Tg-mediated insertion in both *Arabidopsis* and Micro-Tom.

Bayesian optimization of nickel-catalyzed cross-electrophile coupling reactions via high-throughput experimentation

Lukas A. Garcia

Mentor: Sarah E. Reisman

Cross-electrophile couplings (XEC) are a prominent reaction class that have recently been of pharmaceutical interest. Defined as a reductive catalyst-driven bonding of two different electrophiles, C(sp²)-C(sp³) coupling with nickel has been of particular interest, given its ability to bypass previous

shortcomings of typical C-C bond forming reactions, like selectivity and reactivity for one C-H bond. Despite recent advancements, achieving high yielding Ni-catalyzed XEC reactions remain largely difficult due to the large scope of electrophiles, catalysts, and reaction parameters available, making optimization time-consuming and labor intensive. To address these limitations, we apply Bayesian optimization to determine reaction conditions that generalize across a wide range of aryl halide (ArX) and alkyl halide (AlkX) coupling reactions. After five rounds of optimization, we identify conditions that generalize with high yield and reproducibility across the ArX/AlkX scope, using HTE-generated replicate data and average SNR as a generality metric.

Wireless neuroelectronic system for recording and stimulation in freely moving rodents

Said M. Garcia

Mentors: Anqi Zhang and Jialun Zhu

Neurological disorders are the leading cause of disability worldwide. Emerging ultrasoft neural interfaces provide tools for the modulation and monitoring of neural circuits to understand and treat these conditions. To this effort, we developed a wireless embedded system for neural recording and stimulation, designed to interface with implantable electrode arrays in freely moving rodents. The system integrates a Bluetooth Low Energy microcontroller and a 16-channel bidirectional neural recording integrated circuit on a custom 6-layer printed circuit board (22.9 mm × 34.2 mm × 1.6 mm), capable of simultaneous recording and stimulation across 16 of 32 selectable electrodes, with neural data streamed wirelessly to a host computer. A systematic board bring-up protocol was developed and executed to verify each subsystem independently before integration. The device is currently being validated and tested directly on freely moving rodents. This work establishes a wireless hardware and firmware platform for closed-loop neuroelectronic interfaces, with broader implications for minimally invasive treatment of neurological disorders.

Conversational video segmentation

Avni Garg

Mentors: Georgia Gkioxari and Aadarsh Sahoo

Language-grounded video segmentation is essential for spatio-temporal understanding. Referring video object segmentation can spatially localize entities described by simple language expressions and recent work has extended this to reasoning-based prompts. Both settings assume that the target should be segmented throughout the video and emit a mask track per expression with the full clip given offline for the model to see in its entirety before prediction. In this work, we address this gap by introducing Conversational Video Segmentation (CVS) and a benchmark for temporally grounded segmentation, where models must determine not only **what** to segment, but also **when** it should be segmented. We focus on two cases: state conditioned, where the model needs to segment the entity only while it satisfies the prompt (e.g. "the man while he is holding the cup"), and event based prompts, where the model needs to segment the entity before or after the provided event ("the dog before it jumps"). We show that current video segmentation models lack state-aware tracking and are inadequate for CVS. We propose an architecture that can handle this task by predicting a state score for every frame to determine if the entity should be segmented at that time. We also propose a fully automated data generation pipeline for scalable segmentation.

Printed nanowire phototransistors arrays with dual-modulated synaptic plasticity for neuromorphic computing

Ido Geffen

Mentors: Ravinder S. Dahiya and R. Michael Alvarez

Traditional sensing systems typically separate sensing, memory, and computation, requiring large volumes of raw data to be transmitted to centralized processing units. This conventional architecture introduces substantial communication overhead, processing latency, and energy consumption. Neuromorphic devices provide a promising pathway toward low-power edge intelligence by integrating sensing, memory, and information processing directly at the sensor interface. Meanwhile, additive manufacturing and printed electronics can reduce material waste and simplify fabrication processes, offering considerable potential for low-cost, large-area, and scalable device production.

This work reports a fully printed field-effect transistor array based on bottom-up-synthesized ZnO nanowires. An end-to-end fabrication workflow is established, encompassing nanomaterial synthesis, patterned nanowire transfer, device printing, and neuromorphic functional demonstration. Furthermore, the ZnO nanowire transistors respond to programmable ultraviolet light pulse sequences and electrical pulse stimuli, enabling distinct and tunable dual-modal responses. Ultraviolet light pulses are employed to encode visual information, whereas electrical pulses are used to represent tactile inputs. These dual-input characteristics demonstrate the potential of the printed ZnO nanowire transistor array for image recognition and electrically encoded tactile learning, providing a scalable platform for multimodal neuromorphic sensing and edge computing.

How does phenazine retention in *Pseudomonas aeruginosa* biofilms generalize to diverse clinical isolates?

Tasnim Ghoniem

Mentors: Dianne K. Newman, Eloise Masquelier, and Ezekiel J. Johnson

Bacterial biofilms are multicellular communities embedded in a self-produced extracellular matrix. *Pseudomonas aeruginosa* biofilms persist for years within chronically infected hosts, in part because cells in the biofilm core are more tolerant to antibiotics than their actively growing counterparts. Phenazines, small redox-active metabolites, play an essential role in supporting the antibiotic-tolerant cells in the hypoxic core of *P. aeruginosa* biofilms by serving as extracellular electron shuttles. Here, we tested whether the rules governing phenazine retention in the biofilm matrix, previously described in the model strain PA14, generalize to diverse clinical isolates. Four clinical isolates were grown planktonically and as colony biofilms across several media to characterize phenazine retention. We found that pyocyanin (PYO) and phenazine-1-carboxamide (PCN) were consistently retained within these biofilms across strain and media conditions, while phenazine-1-carboxylic acid showed minimal retention. Preliminary results indicate that the amount of eDNA in the matrix dictates the extent of PYO and PCN retention. Thus, the retention mechanism described in the model strain may be a generalizable feature of *P. aeruginosa* biofilm chemistry, positioning chemical matrix-phenazine interactions as a potential angle for disrupting the persistence of antibiotic-tolerant cells in the biofilm core.

Single-shot universal quantum computation with 4D lifted product codes

Ian C. Gill

Mentors: John P. Preskill, Madelyn Cain, and Muzhou Ma

Abstract withheld from publication at mentor's request.

Designing MOF based catalysts through electron density descriptors

Shrish Goel

Mentors: Anima Anandkumar and Danish Khan

Capturing and converting carbon dioxide at scale requires fast, durable and inexpensive catalysts for CO₂ hydration, a reaction nature performs at near diffusion-limited rates with carbonic anhydrase, a zinc enzyme that is too fragile for industrial use. This project aims to design metal-organic framework (MOF) catalysts that mimic the enzyme's active site at the level of its electronic structure. We developed a pipeline that describes a catalytic site with six three-dimensional electron density descriptors, learns a representation of these fields with a Fourier neural operator, and uses a generative model to propose MOF building blocks that are assembled into full structures with PORMAKE. We developed algorithms to extract metal centers and linkers from MOFs and evaluated several methods for computing electron densities. Conventional density functional theory (DFT) codes (ORCA, VASP, and CP2K with density functional embedding) were either inconsistent with one another or too expensive to scale, and the low-cost xTB method produced unreliable reactivity descriptors. We therefore adopted the Kohn-Sham Fourier Neural Operator (KS-FNO), which produced DFT (PBE)-quality densities for a 216-atom MOF in about 16 minutes on a single GPU, and for a representative 150-atom zinc MOF in about 4.5 minutes. Because KS-FNO is currently limited to closed-shell systems, we scoped a pilot dataset of 1,337 closed-shell zinc MOFs below 250 atoms, matching the Zn(II) center of carbonic anhydrase, which covers 88% of node and 98% of edge building-block labels at a

few percent of the full compute cost. We also redesigned descriptor boxes to be centered on metal nodes, which captured the complete coordination environment in every structure tested. These results establish a scalable route to generating the training data needed for electronic mimicry of enzymes in MOFs.

Investigating the characteristics of puncta formed by PPAR γ

Anaya S. Gokarn

Mentors: Shasha Chong and Donghun Lee

Peroxisome proliferator-activated receptor gamma (PPAR γ) is a transcription factor (TF) that regulates expression of genes that affect adipogenesis and lipid metabolism. Previous studies established PPAR γ form puncta at peroxisome proliferator response elements through homotypic interactions of intrinsically disordered regions. PPAR γ is endogenously activated by a fatty acid which promotes its change into an active AF-2 conformation allowing it recruit other coactivators and form a heterodimer with retinoid X receptor to promote transcription. We aim to study how the activation of PPAR γ affects puncta formation. We used rosiglitazone, a strong agonist, to activate PPAR γ in cells. We found that activated wild-type PPAR γ resulted in larger puncta and partition coefficients as compared to inactive PPAR γ . An inactive mutant of PPAR γ , PPAR γ -Y501F, displayed a lower amount of condensate formation pre-activation; once activated, there was a less significant increase in condensate size and partition coefficient as compared to wild-type PPAR γ . These findings show that post-PPAR γ activation, the ligand-responsive AF-2 region contributes to the remodeling of condensates and enlarges them through the recruitment of other PPAR γ molecules and coactivators. Future work can explore the relationship between PPAR γ condensate size and transcription levels.

Growing cells in computers: Surrogate modeling of ion-channel sensitivity in human cortical neurons

Zeynep Goktepe

Mentors: Costas Anastassiou and Nooshin Abdollahi

Detailed neuron models help us understand how ion channels affect a neuron's electrical behavior. However, running these simulations many times can be slow and computationally expensive. In this project, I trained neural networks to predict 14 electrical features from 20 biophysical parameters and the stimulus amplitude. I used simulation data from 48 human cortical neurons representing two cell types. For each cell, I trained separate models for spiking and subthreshold responses. I then used the trained models to measure how sensitive each predicted feature was to changes in each input parameter. Overall, 95.5% of the cell-feature predictions had a validation score of at least 0.5. Subthreshold features were generally predicted more accurately than spiking features. Two independent groups of cells from the same cell type also produced similar parameter rankings and gradient directions. I reported sensitivity using normalized gradients, gradients in physical units, and relative changes called elasticities. These results show that neural-network surrogates can provide fast and reproducible estimates of parameter sensitivity. I am now comparing these results with direct parameter changes in the original neuron simulator.

Model-based reinforcement learning for active flow control

Mahesh Gondi

Mentors: Anima Anandkumar and Jiachen Yao

Active flow control is a challenging research area, and has shown promising advances in the recent past, with applications especially in the aerospace field. The high dimensionality and non-linearity of fluid flows make it particularly difficult to implement classical control methods for the same. Particularly, we focus on a **turbulent channel flow** (TCF) of Reynolds number 180, where we have only **partial observability** (pressure fields at the channel boundary) and it is challenging to capture the underlying fluid dynamics accurately with this limited information. The goal of the control here is to execute **drag reduction** in the channel, and we do this with the blowing and suction of fluid at the boundary walls. In recent times, the application of **Deep Reinforcement Learning** (DRL) as a control agent for these types of environments has shown promising results. Specifically, past work has been established with model-free DRL algorithms for this problem. The current classical benchmark for

drag reduction in TCF environments is **opposition control**, which uses observations from the interior of the channel. Our work aims to first match and then beat this benchmark, using a **model-based** DRL controller. We base our controller on the **TD-MPC2** algorithm, and use a **Neural Operator** based architecture for its world model to better capture the fluid dynamics while having resolution invariance for the actuators and sensors.

Wireless non-LOS BAN channel characterization and communication at mmWave

Sadik O. Görgü

Mentors: Ali Hajimiri and Zachary Szekeres

Increased data requirements in wireless communications have pushed devices towards higher frequencies of operation. One application where this is particularly important is body-area networks (BANs), which enable communication between separate on-body devices. At mmWave frequencies, on-body wireless communication is strongly affected by the body's shadowing and scattering from the surrounding environment, which can significantly reduce link reliability. In this work, we make use of these scattering effects by first characterizing the wireless channel then using the estimated channel response to establish a reliable link between on-body transmitters and receivers. Channel characterization will be performed at millisecond time scales, completing the measurement and data transmission within the channel coherence time. For this purpose, we are designing RF boards operating at 24 GHz using a 1 GHz-bandwidth FMCW chirp to estimate environmental scattering paths and their corresponding phase delays. A 3×3 phased array with digital beamforming adds flexibility to measurements and reduces channel estimation time. Following channel estimation, data transmission will be performed within the 24.0–24.25 GHz ISM band.

Development of a multi-element inductively coupled plasma mass spectrometry method for trace element analysis in human hair

Maria Goullioud

Mentors: François Tissot, Rebecca J. Ryan, and Theo J. Tacail

Hair acts as a biological archive by being an excretory pathway for trace elements found in the human organism and may provide information on long-term exposure. The analysis of a broad range of elements in hair could be used for the screening of numerous diseases, such as cancer. Inductively Coupled Plasma Mass Spectrometry (ICP-MS) is an analytical technique used for quantifying elements in a wide range of sample matrices, including biological samples, such as hair. Trace levels of multiple elements can be measured simultaneously by ICP-MS, however a robust and accurate analytical method for quantifying trace elements in hair must be established. A multi-element calibration strategy and ICP-MS method was developed for the simultaneous analysis of 65 elements in human hair. For the initial analysis, 57 out of 107 analytes measured in in-house hair samples produced valid results according to the established quality control criteria, including limits of detection, limits of quantification, and agreement between isotopes. Future work will focus on modifying and improving the method and calibration strategy and validating the method using the ERM-DB001 human hair reference material.

Constraint of coronal temperature and properties of local low-luminosity active galactic nuclei

Zoe K. Granadoz

Mentors: Fiona A. Harrison and Xiurui Zhao

More than 80% of X-rays in the extragalactic universe are found to be sourced from active galactic nuclei (AGN). The X-ray emission from AGN is thought to originate from a compact, hot corona surrounding the central supermassive black hole. Despite its fundamental role in powering the X-ray emissions, the physical nature of the corona remains poorly understood. I present a study of the physical properties of AGN coronae using a well-defined sample of local, low-luminosity AGN with high-quality X-ray data. By applying state-of-the-art physical models, we obtain robust measurements of key coronal properties, including the electron temperature and optical depth. We compare these

results with those derived for distant, high-luminosity AGN, enabling the first systematic investigation of coronal properties across a broad range of AGN luminosities. I will also discuss the statistical techniques used to derive reliable population constraints and highlight the implications of our results for the physical origin of AGN coronae and for future X-ray observing strategies.

Experimental characterization and validation of the PARSEC Icarus lander primary structure design

Angel O. Guerra

Mentors: Mory Gharib and Jack N. Caldwell

Icarus is PARSEC's vertical takeoff and vertical landing (VTVL) rocket. This project characterized Icarus' primary structure response under vibroacoustic and static loading, refined Icarus' primary structure Geometrical Dimensioning and Tolerancing (GD&T), and benchmarked design choices against NASA/aerospace industry standards. A test article, composed of a bottom section of the Icarus primary structure, was designed and manufactured, and GD&T of Icarus primary structural components was validated during the assembly phase, which prompted redesigns and remanufacturing of necessary components. In the experimental characterization phase, ANSYS simulations of primary structural components in static loading conditions were run and vibroacoustic data was gathered from the test article. Through such testing, design choices can be validated and iterated upon. Further testing using a larger test article incorporating different sections of the Icarus primary structure will better predict the final structure's performance in intense vibrational environments and landing conditions.

Developing an optical spectrometer as a new diagnostic for the Caltech plasma jet experiment

Miguel G. Guerra Zeferino

Mentors: Paul M. Bellan and Joshua Q. Morgan

This project continues the development of a new, fast optical spectrometer for making time-resolved measurements of the plasma in the Bellan Plasma Group's jet experiment at Caltech. During the Caltech jet experiment, a burst of X-rays is observed, which was not initially expected. The timescales on which this happens were too fast for the previously used spectrometer to provide insight into the chain of events occurring when the X-rays are released. The new spectrometer was shown to be at least 8 times faster and aims to help the Caltech researchers understand the mechanisms by which this X-ray burst is produced. During this project, the spectrometer's circuitry was tested and troubleshooting was done until the circuit was stable and functional. The transfer function of the transimpedance amplifiers was measured which showed that the new spectrometer's interframe time can be bounded between 10-125ns. Then, the dispersion of the spectrometer was measured and found to be too low, suggesting that further development is required. Finally, the spectrometer was tested on the plasma. Expectedly, this did not produce meaningful results but a response to the plasma was observed, which is a good proof of concept.

Engineering a cryogenic cold finger for high-frequency characterization of superconducting microwave resonators

Ella Guo

Mentors: Joseph L. Falson and Sandra Glotzer

Quantum phenomena such as superconductivity emerge only under extreme conditions, requiring ultra-low temperatures to reach the superconducting state. In this project, superconducting microwave resonators fabricated from NbO thin films that exhibited magnetic-field-induced superconductivity are being studied beyond conventional direct-current transport measurements to investigate the origin of the superconducting state. These measurements require high-frequency characterization while the resonators are positioned at the center of a superconducting magnet and kept at a temperature near that of a dilution refrigerator, approximately 10 mK. Since the mixing chamber, the coldest stage of the dilution refrigerator, is not located at the center of the superconducting magnet, a thermal connection is needed between the mixing chamber and the resonators. A cryogenic cold finger was therefore designed and is being fabricated to provide this

thermal connection while keeping the resonators positioned at the center of the magnetic field. The primary thermal link is made from high-purity silver, which is annealed to improve its thermal conductivity. Once fabrication is complete, the performance of the cold finger will be tested to determine how effectively it keeps the resonators at ultra-low temperatures during high-frequency measurements in applied magnetic fields.

Infinitely many primes of basic reduction for some abelian sixfolds

Enoch J. Guo

Mentor: Elena Mantovan

Let E be an elliptic curve defined over a number field with at least one real embedding. Elkies proved, by a Euclidstyle argument, that E has supersingular reduction at infinitely many primes. More recently, Li, Mantovan, Pries, and Tang extended this result to certain abelian fourfolds defined over the rationals. In this project, we adapt their methods to further generalize this result to certain genus 6 curves admitting an automorphism of order 7, that is, we prove the Jacobians of such curves have basic reduction at infinitely many primes. More precisely, suppose that C_t is a smooth projective curve with an affine equation of the form $C_t: y^7 = x(x-1)(x-t)$ defined over a number field k . We show that if there exists an embedding σ of k such that $\sigma(J(t)) < 27/4$, where $J(t)$ is the J -function, then the Jacobian of C_t has basic reduction at infinitely many primes. Note that our result again holds for curves defined over arbitrary real number fields (satisfying the embedding condition), in the spirit of Elkies's original proof.

Computing F_K invariants of Montesinos knots via quantum 3j- and 6j-symbols

Weibing Guo

Mentor: Sergei G. Gukov

The F_K invariant is a quantum knot invariant that admits a state-sum description using representations of the quantum group $U_q(\mathfrak{sl}_2)$. This project aims to develop an explicit method for computing the F_K invariant of Montesinos knots by deriving closed-form formulas for the matrices associated with arbitrary rational tangles. Our approach builds on the quantum 3j- and 6j-symbols introduced by Solov'yev and extends the construction beyond the highest weight setting to include mixed highest and lowest weight Verma modules. This requires rewriting the normalized quantum 3j- and 6j-symbols into a more convenient convention, establishing the corresponding graphical identities, and extending the formulas to the negative regime. The resulting framework provides the algebraic foundation for computing F_K invariants of Montesinos knots.

Virtual Merge models in Hopf algebra Markov chains

Nikash Gupta

Mentor: Matilde Marcolli

Merge operations used in models of syntax can be represented as transitions between forests of binary trees. We study Markov chains in which External Merge, virtual Merge, and Sideward Merge are treated as possible basic transitions. The goal is to understand how different choices of basic operations and cost functions affect the long-term behavior of the chain and whether the resulting models remain linguistically meaningful. We examine smaller versions of the state graph, where the adjacency matrices, maximal entropy random walks, and stationary distributions can be computed directly, and work to extend these computations to forests with an arbitrary number of leaves. These examples show that a model using only External Merge and virtual Merge can be periodic, while adding Sideward Merge can introduce self-loops and restore convergence to a stationary distribution. To encourage structure-forming behavior, we consider cost functions depending on the type of Merge operation or on the Shannon entropy of the leaf partition and investigate the resulting asymptotic behavior.

An automated approach to uniform latex film formation on former

Riya Gupta

Mentors: Julia A. Kornfield, Florence Müller, and Siddharth Premnath

Natural rubber is a critical material used across industries, yet global production is heavily dependent on *Hevea brasiliensis*, creating vulnerabilities due to geographic concentration and disease susceptibility. Guayule (*Parthenium argentatum*), a drought-tolerant shrub native to the American Southwest, offers a promising domestic alternative, but requires optimized processing methods for consistent film fabrication. An automated dip-coating system was developed to control immersion, withdrawal, and rotational motion during latex film formation. A three-stage drying process was developed to improve film uniformity on glass slide formers. The parameters of withdrawal speed and angular speed were investigated and optimized based on hevea and guayule latex material properties. Optical Coherence Tomography (OCT) was used to characterize film thickness and uniformity. Future work will optimize withdrawal rates and other processing parameters to improve film uniformity and repeatability. Additional tensile testing will be conducted using smaller specimens to quantify mechanical properties.

A uniformly reduced ultraviolet spectroscopic repository of supernovae from the Hubble Space Telescope

Ishaan Gurazada

Mentors: Mansi M. Kasliwal and Wynn Jacobson-Galan

We present a uniformly reduced spectroscopic repository of all supernova ultraviolet (UV) observations obtained with the Hubble Space Telescope (HST), spanning both the Space Telescope Imaging Spectrograph (STIS) and the Cosmic Origins Spectrograph (COS). Rather than relying on standard pipeline products from the Mikulski Archive for Space Telescopes (MAST), we perform custom reductions for each dataset, including manual trace placement, background subtraction, cosmic-ray rejection, and coadding of split exposures. The resulting spectra cover a wide range of supernova types and epochs, with particular emphasis on early-time observations that capture flash-ionization signatures from circumstellar material shed by the progenitor star in the years before explosion. These high-ionization features, most prominently in the UV, constrain progenitor mass-loss rates and circumstellar material geometry in ways that optical data alone cannot. This repository provides a consistent, calibration-controlled foundation for comparison with non-LTE radiative transfer models and for population-level studies of pre-explosion stellar evolution, and is intended to support both current analyses and future UV missions.

Using conflict graphs to determine the best models against adversarial noise

Amudhan S. Gurumoorthy

Mentors: Daniel Cullina and Konstantin Zuev

Deep learning models are very powerful; they can take objects and words and construct strong relations and classifications, which is very useful for identifying images, recognizing speech, and other scenarios which require greater thought. However, one area where these models often struggle is dealing with adversarial data, noise that is slightly and carefully perturbed to fool a classifier. The level of perturbation that is allowed is known as an adversarial budget, and if classifiers cannot deal with these budgets in a proper manner, then they can be exploited, threatening the stability of models in places like fraud detection and medical imaging. In this project, we analyze many different adversarial budgets at the same time from a theoretical lens using matchings and graph theory, leading to greater robustness on larger samples while keeping smaller budgets to maintain accuracy near a sample point. By representing different samples and different budgets per sample as vertices, we obtain a bipartite conflict graph with special properties that we can exploit. We use extensions and new approaches to algorithms including the Hopcroft-Karp and Hungarian algorithms to calculate and approximate the most number of matchings in this bipartite graph in the least amount of time. This would allow us to determine which adversarial budgets are the best for the classifier to focus on, providing a near-optimal solution to fitting the classifier to be adversarially robust and accurate.

Positroid stratifications of positive symplectic Grassmannians

Henry Gustafson

Mentor: Yassine El Maazouz

Postnikov's positroid decomposition gives a combinatorial cell structure on the totally nonnegative Grassmannian. We study positive symplectic Grassmannians, consisting of nonnegative subspaces that are isotropic for a fixed symplectic form, by intersecting positroid cells with the symplectic equations. For two-planes in six-dimensional space, we classify the nonempty intersections and find six covering relations where the inherited positroid stratification fails. We also study explicit subtraction-free parametrizations of the resulting strata, investigate whether modified plabic graphs or network models can impose the symplectic constraints automatically, and compare the different positive geometries produced by different choices of symplectic form. We also explore extensions from Grassmannians to isotropic partial flag varieties using both positroid coordinates and Marsh–Rietsch parametrizations.

Evaluating causal machine learning algorithms

David Gvazdabia

Mentors: Leonard J. Schulman and Erik Jahn

Causal discovery is the problem of inferring causal relationships between variables from observational data, offering insight beyond correlational or standard machine-learning methods. The problem is typically framed as recovering the directed acyclic graph (DAG) that represents the system, where a directed edge from X to Y indicates that X has a direct causal effect on Y . In recent years, a variety of algorithms have been proposed for this task, but they typically rely on strong assumptions about the data and data-generating process, such as independent and identically distributed (i.i.d.) samples and an acyclic ground-truth graph. These algorithms are also mostly validated on synthetic data, which lacks the complexity and noise of real-world measurements. To better understand how well they actually perform, we introduce computationally efficient baselines that are simplified versions of causal algorithms. We evaluate both the causal algorithms and these baselines on real-world datasets whose data-generating processes are independently known, rather than simulated. Our findings suggest new insights on the challenge of causal discovery on real-world datasets, and that new approaches may be required to achieve better performance.

Spatiotemporal dispersion analysis and theoretical model fitting of self-excited dust acoustic waves in cryogenic water-ice plasma

Muhammed Halil

Mentor: Paul M. Bellan

We investigate the spatiotemporal dispersion characteristics of self-excited dust density waves propagating through dense, levitated cryogenic water-ice grain clouds (Telectrode ≈ 150 K, 13.56 MHz RF power at 40 W) across dynamic neutral gas pressure sweeps. Utilizing high-speed video recording at 1900 frames/s, a time-domain Hilbert transform algorithm is implemented to construct two-dimensional empirical dispersion density maps $S(k, \omega)$ on a fixed 4000×2000 grid matrix (Menzel et al. 2011; Tadsen et al. 2015). Wave frequencies substantially higher than those commonly reported for room-temperature dusty plasmas are observed. Empirical dispersion ridges are systematically evaluated against Thermal DAW, Linear DAW, Discrete Dust Lattice Wave (DLW), and Continuum DLW theoretical models (Rao et al. 1990; Farokhi et al. 1999). We demonstrate that incorporating a low-wavenumber confinement offset ω_0 eliminates systematic residual errors near $k \rightarrow 0$ without requiring unphysical thermal parameter values. More significantly, evaluating the Thermal DAW and Discrete DLW descriptions against the empirical dispersion ridges shows that these two models capture the observed curvature, indicating that the dispersion relation is not strictly linear in wavenumber. Furthermore, dynamic pressure sweeps reveal a clear inverse frequency-pressure relationship ($d\omega/dP < 0$), potentially driven by Epstein neutral gas friction (Flanagan & Goree 2010). The methodology and results presented establish an experimental and analytical frame-work for probing collective modes in cryogenic, non-spherical ice clouds.

Detecting CGM [OII] and [OIII] emission from $z \sim 0.5 - 1.0$ QSOs via stacking

Charis M. Hall

Mentors: Mandy Chen and Charles C. Steidel

Investigating the circumgalactic medium, or CGM, of distant galaxies with actively accreting supermassive black holes (quasars or QSOs) is an important aspect of understanding the nature of galaxy evolution. This is due to the wealth of information the CGM holds about its host galaxy, including contributions to future star formation, dynamics of the galaxy, the fate of mass ejected from the interstellar medium, and the total baryonic mass associated with the galaxy. While CGM emission is intrinsically faint, integral-field spectroscopy (IFS) has enabled direct detection of extended emission, offering a useful way to probe the spatial extent and kinematics of CGM gas. In particular, [OII] and [OIII] emission lines trace ionized, metal-enriched gas and preserve kinematic information that resonantly scattered lines do not, making them well-suited for constraining turbulent velocity fields in the CGM. This project aims to detect [OII] and [OIII] emission around $z \sim 0.5 - 1.0$ QSOs that lack individual prior detections of extended line emission by stacking MUSE IFS data cubes. Stacking increases sensitivity to low surface brightness emission that falls below the detection threshold for single fields. Results from this stacking analysis, including detections or robust upper limits on [OII] and [OIII] emission, will be used to place empirical constraints on the spatial and kinematic structure of ionized CGM gas.

Engineering truly tiny, genetically encodable, acoustic actuators in mammalian cells

Easan S. Hamkins

Mentors: Mikhail G. Shapiro and Anthony Phung

Abstract withheld from publication at mentor's request.

The diameter of random regular graphs

Hypatia C. Hamkins

Mentors: John Urschel and Catherine Babecki

In a graph with a fixed number of vertices, there is a trade-off between the maximum degree of any vertex and the graph's diameter: graphs with smaller diameters invariably have vertices of higher degree. For example, the Moore bound tells us that a 3-regular graph, a graph where each vertex has degree 3, has a diameter of at least $\log_2 n$. The diameter of a random 3-regular graph is known to come close to this lower bound asymptotically, lying between $\lfloor \log_2 n \rfloor + \lfloor \log_2 \log n - \log_2 18 \rfloor + 1$ and $\log_2 n + \log_2 \log n + \log_2 6$. But what exactly the probability distribution of the diameter looks like in this range is an open question. We argue that the diameter of almost every 3-regular random graph is $\text{round}(\log_2 n + \log_2 \log n)$. Under the assumption that there are no dispensable edges (edges that do not result in the discovery of new vertices), we argue that the diameter of a standard configuration model concentrates closely around one value. We then put a tight bound on the total number of dispensable edges that can occur to demonstrate that the diameter still concentrates without this assumption.

Fiber-optic sensing of seismic velocity change in urban settings

Lela R. Hanson

Mentors: Zhongwen Zhan and Elijah Bird

Seismic velocity change is a powerful tool to measure changes in temperature and water content in the subsurface. The ambient seismic wavefield can be used to sample the Earth continuously and at depth, with potential to provide important data for sustainable groundwater management. The technology remains underdeveloped because existing studies of seismic velocity change have been limited by sensor density. In this exploratory study, we leverage Distributed Acoustic Sensing (DAS) to measure surface wave velocity change beneath a 40-kilometer commercial telecommunications cable in Los Angeles at high spatial resolution. We gather pairs of channels in common offset and common midpoint patterns to locate the signal robustly in space and establish sub-seasonal temperature

fluctuations at this site. We also identify the effects of noise source variations on our signal. Our methods and findings offer guidance for other researchers attempting similar methods in urban settings and highlight the ongoing potential of DAS to extend the range and density of ambient seismic field interferometry.

Capturing rare events: A theory-motivated rethinking for attention kernels in transformers

Kailen A. Hargenrader

Mentors: Andrew M. Stuart and Bohan Chen

The transformer architecture underlies frontier models (ChatGPT, Gemini, Claude, etc.). These models have achieved incredible feats in practical tasks such as language modeling. Our work studies a task where transformers can be less successful. The defining feature of the transformer architecture is the attention map. We study the attention map in the distribution mapping setting, where samples from an input distribution are transported through the attention map to become samples from an output distribution. Previous work has shown that transformers are universal operator approximations. Given an operator and an error ϵ , there exists a fixed size transformer that matches the operator to ϵ precision. This approximation considers the attention map as a mapping between distributions, and is agnostic to samples. When considering the attention map as a mapping between samples, the universal operator approximation collapses to a comparable guarantee only for distributions that decay faster than a gaussian. Our work focuses on distributions with heavy tails, where this guarantee does not apply. Heavy-tailed distributions frequently appear in applications such as economics, warranting the study of attention maps between distributions with heavy tails. We devise simple operators on heavy-tailed distributions for which the traditional softmax attention map struggles. Additionally, alternative attention kernels have been proposed in recent literature primarily with the goal of improving computational efficiency. In particular, linear attention variants factor the attention kernel to reduce $O(n^2)$ to $O(n \log n)$ or $O(n)$ runtime. We investigate whether these kernels can outperform softmax in modeling heavy tails.

Growth and characterization of AlF₃ thin films by thermal atomic layer deposition for UV optical applications

Hrishi V. Harsora

Mentors: Austin J. Minnich and Finley B. Donachie

Coronal Mass Ejections (CMEs) are clouds of magnetized plasma ejected from the Sun that can trigger geomagnetic storms, disrupting satellites, communications, and power systems if they reach Earth. The severity of such an event depends on the orientation of the plasma's magnetic field relative to Earth's, which can be inferred from the polarization of the characteristic ultraviolet light emitted by the ejected plasma. Hence, measuring this polarization is key to assessing the threat a CME poses, but it currently relies on bulky instruments spanning several feet. Optical Metasurfaces — flat, nano-chip-scale arrays of precisely shaped nanopillars — could perform this measurement in a compact form, provided the nanostructures are coated with a low-loss dielectric that is transparent in the ultraviolet spectrum and can be applied with nanometre precision. AlF₃ is a wide-bandgap material which meets these requirements. We demonstrate a thermal atomic layer deposition process in a custom-built reactor for the conformal growth of predominantly single-crystalline AlF₃ films on Si and MgO. The films' structure, optical properties, morphology, and composition are characterized by spectroscopic ellipsometry, X-ray diffraction, X-ray reflectivity, atomic force microscopy, and stylus profilometry. This work could enable a new class of compact, lightweight ultraviolet optics for space-based instruments.

The effect of cable bacteria colonization on the marine heat wave tolerance of *Zostera marina*

Quinn J. Hartman

Mentors: Victoria J. Orphan and Magdalene A. MacLean

Marine heat waves are increasingly responsible for large-scale seagrass die-offs worldwide, including a documented loss of eelgrass (*Zostera marina*) at Newport Beach. Heat stress raises sulfide concentrations in marine sediment, and sulfide is toxic to seagrass roots. Cable bacteria, a group of

filamentous bacteria that oxidize sulfide in sediment, may reduce this stress and improve seagrass heat tolerance. This project tests that possibility using a laboratory experiment comparing eelgrass grown in sediment with and without an established cable bacteria population, under both normal and simulated heat wave temperatures. Shoot and root length, sediment sulfide levels, and plant dry mass are being measured before, during, and after the simulated heat wave. Early sulfide measurements show a clear difference between sediment with and without cable bacteria, consistent with cable bacteria actively lowering sulfide levels. Seagrass has not yet shown visible signs of heat stress, but temperature and sulfide measurements are ongoing. Results will help clarify whether cable bacteria populations play a meaningful role in helping seagrass meadows withstand future marine heat waves.

Classical quantum field theory: A field-ontological interacting many-worlds interpretation

Blake W. Hawkins

Mentor: Charles T. Sebens

Quantum theory is our most accurate mathematical framework for describing fundamental physics, but the ontology and structure of physical reality it describes remains a matter of interpretation. In the non-relativistic regime, Newtonian Quantum Mechanics (NQM) offers a metaphysically appealing interacting many-worlds interpretation of quantum mechanics, in which quantum effects arise from interworld interactions between classical point particles. The wavefunction is eliminated from the fundamental ontology and Born-rule probabilities can be naturally understood in terms of self-locating uncertainty. Here, we extend the core principles of NQM to relativistic quantum field theory, constructing a field-ontological interacting many-worlds interpretation in which quantum effects arise from interworld interactions between classical field configurations. Among the implications of CQFT, particular attention is given to the treatment of energy.

Evaluating causal machine learning algorithms

Alexander He

Mentors: Leonard J. Schulman and Erik Jahn

Causal discovery is the problem of inferring causal relationships between variables from observational data, offering insight beyond correlational or standard machine-learning methods. The problem is typically framed as recovering the directed acyclic graph (DAG) that represents the system, where a directed edge from X to Y indicates that X has a direct causal effect on Y . In recent years, a variety of algorithms have been proposed for this task, but they typically rely on strong assumptions about the data and data-generating process, such as independent and identically distributed (i.i.d.) samples and an acyclic ground-truth graph. These algorithms are also mostly validated on synthetic data, which lacks the complexity and noise of real-world measurements. To better understand how well they actually perform, we introduce computationally efficient baselines that are simplified versions of causal algorithms. We evaluate both the causal algorithms and these baselines on real-world datasets whose data-generating processes are independently known, rather than simulated. Our findings suggest new insights on the challenge of causal discovery on real-world datasets, and that new approaches may be required to achieve better performance.

Density degree set of Algebraic Varieties

Xi He

Mentor: Yu Fu

We study density degree sets of algebraic varieties over number fields, extending the sweeping-curve method for products of two curves to higher-dimensional products of curves and to products with Severi–Brauer varieties.

For a product $V = C_1 \times \cdots \times C_n$ of nice curves, each with a k -rational point, we construct Zariski-dense families of iterated fiber-product curves. We obtain an explicit effective threshold, in terms of the genera and gonality of the factors, beyond which every degree belongs to $\delta(V/k)$. For products of positive-genus curves without rational-points, we establish a sharper zero-cycle-sensitive bound and deduce a uniform bound in terms of the genera and effective index.

Let C be a nice curve, let A be a central simple k -algebra of index r , and let $S = \text{SB}(A)$. For every degree- r splitting field K/k of A , we prove

$$r \delta(C_K/K) \subseteq \delta(C \times S/k) \subseteq \delta(C/k) \cap r\mathbb{N}.$$

For an elliptic curve E and nonsplit A , we further show that $r \in \delta(E \times S/k)$ if and only if E has positive Mordell–Weil rank over some degree- r splitting field of A . We also prove $\wp(C \times S/k) = \wp(C/k)$, where $\wp$ denotes the potential density degree set.

As low-degree applications, we give a branch-locus criterion implying $7 \in \delta(E \times C/k)$ for an elliptic curve E and a nice genus-2 curve C with a k -rational non-Weierstrass point. We also generalize a local obstruction criterion for non-dense quadratic points on products of genus-2 curves from $p = 3$ to arbitrary odd primes, systematically searching new explicit examples.

Investigating groundwater-driven land deformation in the San Gabriel Valley using InSAR time series analysis

Natalia M. Hernandez

Mentor: Mark Simons

Groundwater withdrawals can cause aquifer compaction and land subsidence in groundwater basins such as the San Gabriel Valley. This project investigates the spatial and temporal patterns of land-surface deformation across the San Gabriel Valley and their relationship with groundwater variability. Interferometric synthetic aperture radar uses repeat-pass radar images to detect and map changes in the land surface over time. Satellite observations from Sentinel-1 are then processed to generate displacement time series that show how the ground has moved throughout the years of 2014-2026. These time series are then analyzed using mathematical models that capture the simplest temporal behavior supported by the observations, including long-term seasonal deformation where appropriate. Groundwater level extraction records are then compared with the observed deformation to evaluate whether changes in groundwater conditions respond to subsidence or uplift. Continued analysis will determine which temporal models best describe deformation in the San Gabriel Valley and help assess how this complicated groundwater system responds to changes in subsidence and storage

Accelerating a closed-loop design-build-test pipeline for AI-designed proteins with multiplexed Sidewinder DNA construction

William Cheung Ho

Mentor: Kaihang Wang

Artificial Intelligence has enabled the guided exploration of the infinite protein sequence space at unprecedented speed and scale. The rise of AI-driven protein design has outpaced the capacity to construct DNA at scale, leaving a surplus of in silico sequences unrealized in the physical world. Here we apply Oligo Pool Sidewinder assembly, which relies on Sidewinder's DNA three-way junction to construct DNA accurately and efficiently. Using Foldtuning, we constructed 180 AI-designed sequences predicted to adopt the structure of GFP across six oligo pools (thirty sequences per pool). Taking the narrow fitness landscape of GFP into consideration, we optimized an *E. coli* expression system to aim for the detection of candidates with low fluorescence intensity. Ultimately, we underscore the limits of in silico structural prediction, demonstrating the need for physical realization to improve protein models.

Streamlining the copper gasket making process

Owen M. Hocks

Mentors: Richard McCarthy, Jon Feicht, Jordan Vanosky, and Melina Fuentes-Garcia

Because buying from local manufacturers became inconsistent and buying from overseas became too expensive, the solution shifted to making copper gaskets in-house from a spool of 99.99% copper. Where the copper weir is then straightened, roughly measured to the desired length, cold welded, then grinded/polished for 20 minutes or more, using a Dremel. Thus resulting in about 25-30 minutes to make one gasket. After the gasket is made it's then sealed, pumped down, and leak-checked. But if

the leak-check fails the gasket needs to be replaced due to it being a one time use. So due to the measuring and grinding/polishing system taking the majority of the time that is where the attention is focused. Making the newer system easy to operate and less taxing on the whomever is making the gasket. Getting the time down to 5-10 minutes for one gasket, and creating more time for the user to get other tasks done during their day. This also creates a consistency between each of the copper gaskets, saving the need to redo the process when a gasket fails.

Physical plausibility emerges from resilience across diverse noise paths in video diffusion models

Richard Hoffmann

Mentors: Yilun Du, Georgia Gkioxari, Ruojin Cai, and Sarah Liaw

Pre-trained video diffusion models excel at generating visually realistic rollouts, but their predictions often lack an understanding of real-world dynamics. This can lead to physically implausible state transitions, temporal artifacts, and unexpected hallucinations. However, generated videos can still simulate future objections in robot planning and decision-making settings, making the use of video generation models as world models for embodied agents a dominant paradigm. For generalizable robot learning beyond controlled settings, it is imperative that video models better represent real-world dynamics. Previous work improve generated rollouts by sampling more trajectories to increase the chance of outputting a coherent rollout, but standard best-of-K methods require an external verifier, hand-tuned rewards, and inefficiently delete information from unsuccessful samples. Instead of relying on retraining or additional external verification methods, we propose an inference-time debate algorithm that treats independently initialized denoising trajectories as individual agents proposing alternate futures. Agents challenge proposals by re-noising with their own initial noise tensors and querying the video model to reconstruct it. We claim that a more physically plausible rollout is one that can be recovered across several counterfactual noise directions, which we evaluate on the Physics-IQ benchmark.

Flume experiments on riverbank erosion with smooth mud walls

Natalie R. Homyk

Mentors: Michael P. Lamb, Kimberly L. Miller, and Maria Schmeer

Bank erodibility sets river channel geometry and bank migration rates, directly shaping floodplains, human infrastructure, and local ecosystems. Natural rivers often exhibit boundary stresses far exceeding the entrainment thresholds of sand or gravel, supporting the hypothesis that cohesive mud supports bank structure. However, testing these mechanics in natural channels is challenging due to complexities and spatial variation, and controlled flume experiments directly evaluating bank erosion in cohesive mud are lacking. To address this gap, we performed open-channel flume experiments testing riverbank erosion theories on cohesive mud walls across varied boundary shear stresses.

In our experiments, bank retreat did not occur through individual grain entrainment, but rather through localized cracking, block detachment, and slumping. Additionally, observed erosion thresholds and rates conform to classic shear-stress theory only when applying an erosion threshold sufficiently high for cohesive mud and explicitly correcting boundary shear stresses for wall smoothness. Notably, this boundary stress partition reduces the wall shear stress relative to the channel bed by an order of magnitude. Our results indicate that applying threshold channel theory without accounting for hydraulic stress partitioning on smooth, cohesive banks may lead to substantial overestimates of bank shear stress and bank strength in natural rivers.

Using visual salience manipulation to promote healthier food choices in vending environments

Alexa M. Hong

Mentors: Colin F. Camerer and Shiang-Hung Hu

Ultra-processed foods (UPFs) are prevalent in modern food environments and contribute to poor dietary quality and adverse health outcomes. While many interventions encourage healthier eating through pricing, labeling, or product placement, fewer target the perceptual mechanisms that guide

attention before conscious decision-making. This project investigates whether manipulating bottom-up visual salience can influence consumers to choose less ultra-processed foods in vending machine environments without restricting choice. Using a computational salience algorithm (Saliency Attentive Model, SAM), images of food and beverage items will be modified to create *salience makeovers* that increase attention to lower-UPF products and *salience makeunders* that reduce attention to higher-UPF products while preserving product recognizability. The effectiveness of these manipulations will be evaluated in two complementary experiments: a campus vending machine field study using item-level sales data and a controlled online vending machine simulation in which participants make hypothetical purchase decisions. Across both settings, product price, availability, and choice structure will remain constant, isolating the effects of visual salience on consumer behavior. The study is expected to determine whether algorithmically manipulating perceptual attention can reliably shift food choices toward healthier alternatives. By integrating computational vision, cognitive psychology, and behavioral economics, this research advances understanding of how bottom-up attentional processes shape decision-making and explores a scalable, low-cost intervention for promoting healthier eating in real-world choice environments.

Learning exact exchange for Kohn-Sham density functional theory

Arimitsu Horikawa-Strakovsky

Mentors: Anima Anandkumar, Danish Khan, and Maurice Hanisch

Kohn-Sham density functional theory (DFT) is one of the main methods utilized in modern quantum chemistry for electronic structure calculations. Semi-local functionals such as PBE suffer from self-interaction error, leading to systematic errors in calculations of many properties such as reaction barrier heights. Hybrid functionals such as PBE0 mitigate this issue by incorporating a fraction of exact exchange, substantially improving accuracy, but at the cost of increasing computational scaling from $O(N^3)$ to $O(N^4)$. Thus, we learn exact exchange using a continuous radial factored Fourier neural operator (cRFNO). Training on only 500 neutral molecules, we are able to approach the accuracy of hybrid functionals for reaction barrier heights on BH76, achieving errors within 2.2 kcal/mol of PBE0 compared to PBE's error of 5 kcal/mol, while retaining the scaling of semi-local DFT. Despite training only on neutral molecules, our model transfers well to charged molecules and reactions, demonstrating strong generalization. We also observe substantial improvements over PBE in electron density and total energy, halving the error relative to PBE0. These results suggest that our approach offers a promising path toward emulating the accuracy of hybrid functionals at a fraction of their computational cost.

Efficiently steering diffusion models with a coarse reference

Max M. Horwitz

Mentors: Adam C. Wierman, Lillian Ratliff, and Yashaswini Murthy

Diffusion models are capable of learning complex multimodal distributions over actions, images and configurations. When deploying these models in practice, one often wishes to combine a powerful generalist model with some task-specific information. Most existing algorithms require additional training, either fine tuning the base model or learning a noise policy over the latent space. This additional training can be expensive, which has motivated an alternative approach: using a coarse reference to steer the diffusion process. In this work, we present a theoretical analysis of the coarse steering framework and a recently proposed algorithm (Flow Reversal Steering). While originally proposed for robotics, this framework has the potential to be useful in other domains such as image reconstruction and generative discovery. We also study new algorithms for this setting, using tools from optimal control, and present experimental results.

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.

Function-space certification for operator learning

Ryan L. Hsiang

Mentors: Anima Anandkumar and Robert Joseph George

Certification for scientific machine-learning models is essential when robustness and physical consistency must hold beyond a finite test set. Existing neural-network verification methods are designed for finite-dimensional vectors. When applied to a discretized neural operator, they certify only the selected grid points and provide no guarantees over the domain or across another resolution. We close this gap by formulating interval bound propagation (IBP) and CROWN directly in function space. We propagate rigorous functional bounds with validated integration using numerical quadrature with Lipschitz-based error corrections and arbitrary-precision ball arithmetic. We experiment on Fourier neural operators trained for Poisson, Helmholtz, and Burgers equations. Against Monte Carlo sampling and projected-gradient attacks, function-space CROWN produces bounds within $1.35\sim 4.55\times$ of the empirical attacks across the three PDEs, providing tight, certified enclosure over the entire continuous domain.

Probing atmospheric mass loss from WASP-33 b with metastable helium

Cheng-an Hsieh

Mentors: Heather A. Knutson, Morgan Sidel, and Pablo Drake

A stars have high-energy spectra fundamentally distinct from those of later-type stars, but their effects on planetary atmospheric mass loss remain unclear. Only two A-star gas giants currently have published mass-loss measurements, both from Balmer-line absorption. Here, we investigate He I 10830 Å absorption in one of these planets, the ultra-hot Jupiter WASP-33b. We observed one transit with Palomar/WIRC and jointly modeled the transit and the host star's δ Scuti variability. We measured a narrowband helium excess of $\Delta\delta\text{He} = 2395(+863/-849)$ ppm (2.8σ), which we fit with sunbather Parker-Cloudy models. Matching this depth requires a relatively high stellar XUV flux, corresponding to $\log_{10}(\dot{M} / \text{g s}^{-1}) = 12.00(+0.66/-0.69)$, consistent with the published Balmer-line estimate. We then compared the observed Balmer and He I absorption with predictions from self-consistent pyTPCI models and explored how their relative strengths vary with planet mass and stellar XUV spectrum. We also present new 3D outflow models that better reproduce both line strengths simultaneously. These results show how joint hydrogen and helium absorption measurements can constrain the physical structure of escaping atmospheres around close-in planets orbiting A stars.

PosteriorBench: From point estimates to posterior matching in evaluating generative inverse solvers

Zi-Siang Hsu

Mentors: Anima Anandkumar and Jiachen Yao

Generative models are increasingly used to solve scientific inverse problems, but existing evaluations still focus primarily on whether a method can produce a single plausible reconstruction. This is insufficient for ill-posed problems, where multiple solutions may be consistent with the same sparse or

noisy observations. In these settings, a method can achieve strong pointwise accuracy while still failing to capture the true posterior through mode collapse, overconfident uncertainty, or averaging incompatible solutions.

We introduce PosteriorBench, a benchmark for evaluating the distributional accuracy of generative inverse solvers. PosteriorBench evaluates four physics-based inverse problems: Darcy flow inversion, Poisson source recovery, carbon capture and storage, and light transport material inference. For each task, we construct high-fidelity reference posteriors using computationally heavy but established procedures such as rejection sampling and Markov chain Monte Carlo, enabling direct assessment of whether solvers recover the full set of solutions rather than the single best sample.

The benchmark spans sparse sensing, low-resolution observations, nonlinear forward models, varying noise levels, and multimodal priors, with a unified pipeline for distribution matching, uncertainty quantification, and hyperparameter calibration. Our experiments reveal substantial distribution-matching gaps across current solvers, while showing that neural operators improve resolution robustness and that guidance weights and generation noise are key to posterior-variance calibration.

Developing a charged palladium catalyst to synthesize photodegradable polyethylene

Chloe I. Hu Pesh

Mentors: Theodor Agapie and Jocelyn Y. Zhang

Polyethylene persists in the environment long after its disposal because its inert carbon-carbon bonds are not readily hydrolyzed and resist photodegradation. When the polymer does eventually degrade, this degradation begins at sites of defects including unsaturated C=C bonds and carbonyls, making the intentional incorporation of these functionalities to the polymer an area of interest. Ethylene/carbon monoxide (E/CO) copolymers which include ketones have been found to be photodegradable while maintaining desirable commercial properties. Although researchers have only achieved single CO insertion in polyethylene, back-to-back CO insertion, which would create an even more photodegradable polymer has not been observed. This project works towards the synthesis of a charged palladium catalyst for the back-to-back insertion of CO into E/CO copolymers. The synthesized anionic catalyst will carry cationic counterions, which we hypothesize will form a thermodynamically favored 5-membered chelate with the diketone functionality formed after the second CO insertion, making back-to-back CO insertion more favorable.

An iodine-referenced laser stabilization system for high-precision cavity ringdown spectroscopy of electron-enhanced atomic layer processes (EE-ALP)

Annie M. Hu

Mentors: Austin J. Minnich and Emily Shi

Electron-enhanced atomic layer processing (EE-ALP) uses electrons as precursors to stimulate and enhance surface chemical reactions during thin film growth. EE-ALP opens new process windows that may help to overcome the limitations of existing ALP methods. However, a fundamental understanding of the chemical mechanisms that occur during EE-ALP is lacking. This project combines a HCP electron source with an in-situ cavity ringdown spectroscopy (CRDS) system in a custom-built reaction chamber, aiming to enable characterization of gas-phase and surface chemistry during growth. One of the preliminary goals for our system was to use CRDS to observe how electron irradiation changes the radical and ion content of common ALP chamber background gases, such as methane. CRDS is an extremely sensitive technique that measures the time constant of decay for built-up laser intensity within an optical cavity. A requirement for precise CRDS is a cavity locked to a frequency-stabilized laser; stabilization ensures the laser remains resonant with the cavity in spite of environmental disturbance, giving reproducible ringdown measurements. Typically, this is achieved by referencing the laser to fine rovibrational molecular transitions. In our setup we work towards stabilizing the beam to a molecular I₂ peak, frequency-modulating the beam as it passes through an I₂ cell. We then detect the discriminator signal of the absorption peak, or of a finer feature known as a lamb dip (on the order

of MHz wide), either of which we can lock to using a servo controller. Experimentation with the optical configuration and modulation parameters is ongoing in order to achieve better signal-to-noise and optimize our lock. Finally, the Pound-Drever-Hall technique is used to lock cavity length to our stabilized laser beam.

Characterization and optimization of Kinetic Inductance Phonon-Mediated (KIPM) detectors for low-threshold dark matter direct detection

Jonathan X. Hu

Mentor: Sunil Golwala

Current dark matter searches have focused on the direct detection of Weakly Interacting Massive Particles (WIMPs) with masses above 1 GeV, whose expected recoil energies lie in the keV regime. These experiments have reached unprecedented sensitivities in the high mass range but have not yet observed WIMP signals, motivating a shift of attention toward sub-GeV dark matter candidates. Phonon-based detectors are a promising technology for resolving the eV and sub-eV energy depositions expected from low-mass dark matter interactions. Kinetic Inductance Phonon-Mediated (KIPM) detectors use superconducting microwave resonators coupled to a target substrate to sense the phonons generated by energy depositions within it. However, the energy resolution of current KIPM detectors is limited primarily by amplifier noise in the readout chain. Kinetic Inductance Traveling Wave Parametric Amplifiers (KI-TWPAs) exploit the nonlinear kinetic inductance of a superconducting transmission line to transfer power from a strong pump tone to a weak signal tone, amplifying the signal parametrically. Their high gain and near quantum-limited added noise can therefore be used to suppress the readout noise and improve the energy resolution of KIPM detectors. In this work, we set up a new cryogenic testing environment in which KIPM detectors and KI-TWPAs can be operated, and we measure the gain and noise of the KI-TWPA. In addition, current models of KIPM detector pulses are largely phenomenological. We therefore work to develop physical models of these pulses based on phonon recycling and quasiparticle pair-breaking mechanisms in order to better understand the detector response. Together, these efforts can help guide improvements in the energy resolution and performance of future KIPM detectors for low-threshold dark matter detection.

SINDy-LLM: Large language model-augmented framework for the sparse identification of nonlinear dynamics

Evan T. Huang

Mentors: Steven Brunton, Adam C. Wierman, and Michael Garcia

Obtaining the equations that model physical phenomena has long been an important but difficult task in science and engineering. Sparse Identification of Nonlinear Dynamics (SINDy) is a powerful data-driven method for discovering these governing equations; however, effective use often requires deep domain knowledge and coordination across specialized tools. In this project, we introduce SINDy-LLM, an agentic, tool-augmented framework for LLM-assisted equation discovery. SINDy-LLM combines callable scientific tools—SINDy variants, denoising, skeleton fitting, and other algorithms—with Retrieval-Augmented Generation (RAG) and a QLoRA (Quantized Low-Rank Adaptation) fine-tuning pipeline for teaching the model scientific workflow behavior. It is evaluated on both canonical and nonlinear benchmark data systems using symbolic accuracy, term-level F1 (term accuracy), and normalized mean squared error (NMSE).

Reducing error rates in *de novo* DNA synthesis

Jerry P. Huang

Mentors: Kaihang Wang and Daniel Graves

De novo DNA synthesis technology is at the cornerstone of many modern biological advances and fields: synthetic biology, genetic engineering, and personalized medicines. Despite its importance, *de novo* DNA synthesis has a relatively high error rate (10^{-3} to 10^{-4}) due to the cyclical and error-compounding nature of phosphoramidite synthesis. Here, we focus on developing filtering technology that can catch and remove mutations. One approach, which we call “basewise correction,” provides DNA polymerase with only the correct next base, one at a time, such that the polymerase can only replicate the correct sequence (excluding truncation mutations). Work has been done to improve the

cyclical yield rate of basewise correction from 4% to 94%, yet further work is needed to maintain these high yield rates over many cycles. Another approach is enzymatic filtration, in which *in vivo* DNA repair enzymes such as EndoMS are repurposed *in vitro*. However, EndoMS has bias for certain types of mismatches (specifically, G/G, G/T, and T/T), and cannot detect others, such as A/A. LigandMPNN was used to redesign the DNA mismatch binding pocket of EndoMS to fit A/A, and future work will be directed at validating the redesigned proteins.

Ultrasound-induced flow control to reduce bacterial migration in urinary catheters

Ting Syuan Huang

Mentors: Chiara Daraio and Amanda Y. Lin

Urinary catheters can provide a pathway for bacteria to move toward the bladder and form biofilms that contribute to infection and blockage. Existing prevention strategies, including antimicrobial coatings, passive surface modifications, and active cleaning systems, often suffer from limited long-term effectiveness. In this project, we investigate whether ultrasound can manipulate the flow field inside the catheter to suppress bacterial migration along the catheter wall without altering the catheter surface. A two-dimensional catheter model was developed in COMSOL Multiphysics to simulate acoustically induced flow generated by an ultrasonic transducer. The resulting flow fields were coupled with a bacterial dynamics model to predict the spatial distribution of bacteria over time. The simulations demonstrate that ultrasound significantly modifies the near-wall flow and reduces upstream bacteria migration. By systematically varying the transducer velocity and background flow speed, we identified operating conditions that maximize downstream bacterial transport. These findings suggest that ultrasound-induced flow manipulation is a promising strategy for reducing bacterial migration and preventing biofilm formation in urinary catheters.

HomeBody: A humanoid that explores, remembers, and acts on its own

Gio Huh

Mentors: Karen Liu, Georgia Gkioxari, Guy Tevet, and Takara Truong

A hierarchical approach to humanoid loco-manipulation connects task reasoning (System 2) to motor execution (System 0) through a learned action policy (System 1). This intermediate policy converts visual context and instructions into motion targets. As frontier vision-language models become more capable, we revisit whether System 2 can directly select reusable skills and spatial targets for System 0. We build HomeBody to demonstrate the feasibility of this interface for long-horizon humanoid loco-manipulation. To identify where to act, HomeBody builds persistent spatial memory from the humanoid's own exploration, linking camera observations to measured room geometry. System 2 uses the spatial memory to select a skill and target, while the skill uses local perception and control to execute the motion. On a Unitree G1, this loop supports room-wide tidying and retrieval of medicine from a remembered drawer in a previously unseen kitchen. These qualitative demonstrations use pretrained components without additional environment-specific policy training or task demonstrations.

Investigating self-replicating droplets

Gene G. Huntley

Mentors: Erik Winfree, Cameron Chalk, and Yancheng Du

Programmable DNA nanostar condensates provide a model system for studying how molecular interactions control phase separation. Motivated by the broader goal of developing self-replicating condensates, we investigated how different types of linker-mediated interactions can produce temperature-dependent changes in phase behavior. Linker strands, previously used to mediate interactions between otherwise non-interacting DNA-functionalized structures, can be designed so that temperature changes alter their binding configurations and effective interaction strengths. We developed a mean-field free-energy framework to characterize how nanostar concentration, linker availability, and binding energetics influence the stability of mixed and demixed phases. Guided by this analysis, we explored temperature-responsive linker nanostructures and computationally screened DNA sequences whose binding energetics vary over the desired temperature range. Together, these results provide a theoretical and design framework for controlling phase behavior in DNA nanostar condensates through programmable linker interactions.

Proteomic analysis of NAC interactions

Lyna T. Huynh

Mentors: Shu-ou Shan, Radoslaw J. Gora, and Sunny Huang

The nascent polypeptide-associated complex (NAC) coordinates co-translational protein biogenesis at the ribosomal exit tunnel, yet the domain-specific interactomes of its flexible α -N, α -C, and β -C termini remain incompletely defined. Here, we employ TurboID-based proximity labeling coupled with quantitative LC-MS/MS proteomics to systematically map terminal-specific NAC interactions in HEK293T cells. In preliminary single-run data, known co-translational factors (including Naa acetyltransferases, NMT1/2, METAP1/2, SRP, and HYPK1) showed enrichment at NAC termini relative to TurboID-only controls, validating that our proximity labeling assay captures specific interactions. Several uncharacterized candidates, including taxilin family members, prefoldin subunits, and annexins, also met our preliminary enrichment cutoff and await statistical validation from ongoing triplicate proteomic analysis. In parallel, we are optimizing an orthogonal RNA Affinity Purification with Poly-Lysine (RAPPL) assay to validate candidate interactions under conditions that preserve ribosome-associated interactions. Together, this work establishes a robust pipeline for quantitative interactome analysis and provides a foundation for defining how individual NAC termini coordinate the recruitment of protein biogenesis factors at the ribosomal exit tunnel.

Nanobody-mediated stabilization of human effectors enables antigen-dependent cell selection

Noor E. Ibrahim

Mentors: Katie Galloway and Rob B. Phillips

Isolating a target population of cells from a mixture based on an intracellular molecular marker, also known as selection, is a central challenge in cell engineering. DASIT (Destabilized-nanobody Antigen Selection and Identification Tool) achieves this by converting an intracellular antigen into a selectable marker. A nanobody carrying destabilizing framework mutations is degraded by the proteasome unless it binds to its target. Fusing it to an effector confers the same stability properties to the fusion protein, resulting in the effector accumulating only in antigen-positive cells. This enriches a population based on intrinsic cell state rather than delivery. Because the mechanism operates at the protein level, it needs no permanent genetic integration and supports transient recyclable delivery. DASIT's central limitation for translational applications is that it has only been shown to function with bacterial resistance markers, and foreign proteins provoke the rejection of the cells that are being selected. Here, we humanize DASIT by replacing the bacterial effector with a panel of human effector proteins that each turn antigen presence into a survival advantage through a distinct pathway, from suppressing stress-induced death with a dominant-negative p53 fragment to upregulating proliferative signaling through cytokine-receptor and STAT modules. Because these effectors function through different pathways, they are used to test whether antigen-based selection generalizes across effectors. The outlined validation spans immortalized lines to iPSCs, with primary T cells as the therapeutic endpoint. Together, this establishes antigen-dependent selection from clinically compatible parts, positioning the effector as a tunable input for selecting and enriching human cells on intracellular state.

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.

Developing a disordered kinetics framework to link soil carbon age, depth, and temperature sensitivity

Felix X. Jacobs

Mentors: Christian Frankenberg and Eva-Marie Metz

Soils store more organic carbon than the atmosphere and global vegetation combined, with changes in stocks reflecting the balance between primary productivity inputs and respiratory losses. The processes underpinning soil carbon turnover remain poorly constrained and are a major source of uncertainty in carbon-cycle projections. Current-generation Earth system models simulate turnover too rapidly, underestimating the mean age of soil carbon several-fold relative to radiocarbon measurements. Better constraining microbial decomposition and its temperature sensitivity is therefore important for quantifying climate-carbon feedbacks.

Here we represent the activation energies of soil organic carbon as a Gaussian distribution which, through an Arrhenius relationship, yields a lognormal spectrum of decay rates and links intrinsic carbon reactivity to its temperature sensitivity. When coupled to vertical transport and seasonal temperature forcing, the model captures broad radiocarbon patterns across climate and depth, with older carbon at depth arising from the preferential survival of slowly decomposing carbon during transport. In idealized warming simulations, an early rise in respiration precedes carbon losses over decades to millennia, with an increasing contribution from older, more temperature-sensitive carbon. The framework is formulated as standard compartmental ordinary differential equations, offering a parsimonious representation of the coupled patterns of soil carbon age and reactivity that remain challenging for Earth system models.

Modelling the effects of tungsten alloying and surface oxidation on coverage-dependent hydrogen evolution in Ni_{17}W_3

Aditya Jain

Mentors: William A. Goddard III and Sejun Kim

Developing a stable, efficient and affordable catalyst for the hydrogen evolution reaction is crucial for cheap renewable energy storage solutions, and nickel-tungsten alloys are among the most promising alternatives to platinum for it. A multicomponent $\text{Ni}_{17}\text{W}_3/\text{WO}_{3-x}/\text{MoO}_{3-x}$ catalyst was shown to reach 10 mA cm^{-2} at an overpotential of 14 mV. In practice the working surface is neither clean nor sparsely covered: it carries close to a monolayer of hydrogen, and X-ray photoelectron spectra of this material show tungsten in partially oxidised states alongside the metal. Here we use density functional theory to build hydrogen coverage up site by site on pure Ni, on Ni_{17}W_3 , and on Ni_{17}W_3 with oxidised surface tungsten. Because each surface becomes thermoneutral at its own coverage, what distinguishes them is how much hydrogen each holds: the oxidised alloy saturates at 0.79 monolayers, while pure nickel is still adsorbing hydrogen at a full monolayer. Searching for hydrogen recombination across this whole coverage range, we find no pathway fast enough to carry the measured current. Hydrogen must therefore be released either by the electrochemical Heyrovsky step, or with the help of the oxide support that our metal-only model leaves out, and pursuing these two routes is the next test of how this catalyst works.

Mapping the inner geometry of active galactic nuclei using Wiener Deconvolution of X-ray reverberation

Adam M. Janik

Mentors: Fiona A. Harrison and Demet Kirmizibayrak

Observations of X-ray reverberation time-lags provide a powerful probe of the geometry of space surrounding supermassive black holes in the center of active galactic nuclei (AGN). Current techniques for measuring these lags primarily rely on Fourier-domain cross-correlation analyses. While effective, these methods recover only aggregate time-lag properties of the system rather than directly resolving the impulse response function that encodes its underlying geometry.

The novel astronomical Wiener Deconvolution (WD) technique not only improves the accuracy and precision of time-lag measurements by suppressing instrumental noise, but also directly estimates the response function, which holds physical information about the system's geometry and light-travel delay processes. In this work, we apply the WD method to AGN with previously detected X-ray reverberation lags. By reconstructing the response function and searching for lags, we investigate whether Wiener deconvolution can place tighter constraints on the geometry of the innermost regions surrounding the central black holes.

Identifying biomarkers for insect herbivory through stable isotope analysis

Bianca D. Jaramillo

Mentor: Julia Tejada

The evolution of herbivory is one of the most consequential ecological transitions in the history of life. By enabling insects to exploit plants as a food source, it fueled a reciprocal evolutionary arms race with flowering plants that drove extraordinary diversification in both groups and helped shape modern terrestrial biodiversity. Yet despite its evolutionary importance, we still know surprisingly little about the physiological innovations that repeatedly enabled insects to overcome plant defenses and evolve herbivory. I propose to address this question by using compound-specific stable isotope analysis of amino acids to identify biochemical signatures of different herbivorous feeding strategies, providing a physiological perspective that complements and extends beyond traditional morphological approaches. I will then integrate these data with published mammalian and arthropod datasets to determine whether amino acid isotopic signatures of herbivory reflect conserved metabolic processes across animal phyla or lineage-specific evolutionary adaptations. Together, this work will establish a new biogeochemical framework for studying the evolution of herbivory in living insects while laying the foundation for reconstructing plant-insect interactions in the fossil record.

Cooling a many-body spin system in diamond using optically reset NV centers

Thuwaragesh Jayachandran

Mentors: Norman Yao, Andrei Faraon, Che Liu, and Haopu Yang

A major limitation of solid state quantum simulation platforms is that an interacting many-body system generally cannot be initialized in a low energy state. The natural starting point is an effectively infinite temperature state. We demonstrate cooling of an interacting ensemble of P1 centers in diamond, a many-body spin system, using nitrogen vacancy (NV) centers as a controllable bath. The NV is dressed by a spin-locking type microwave drive which tunes its effective energy splitting into resonance with the few MHz P1–P1 dipolar interaction scale. Resonant flip-flops then transfer energy from the P1 ensemble to the NV, and optical reset of the NV carries that energy out of the system. The cooling effect is probed by the difference in the spin-locking relaxation time scales which are limited by the noise generated from the P1 ensemble at different cooled temperatures. These results establish bath based cooling as a practical route to preparing low entropy states in interacting dipolar spin systems.

Strategic optimism in games

Oscar Jenner

Mentors: Eric V. Mazumdar, Boris Velasevic, and Nicolas Lanzetti

Strategic optimism, that is, optimism about the other players' strategies, can be beneficial in coordination games, where such optimism is natural. In the spirit of recent work on strategic risk aversion in games, this report studies strategic optimism in finite games and finite-horizon Markov games, focusing on when the resulting equilibria are unique and tractable to compute. We introduce the strategically optimistic quantal equilibrium, which combines optimism with bounded rationality, and show that it exists in every finite game. Independently of the game payoffs, we give conditions on the players' levels of optimism and bounded rationality that guarantee uniqueness of the equilibrium, as well as linear convergence of best responses to it. These guarantees transfer to finite-horizon Markov games, where the equilibrium is computed through backward induction, and the equilibrium remains an approximate equilibrium when the transition kernels are estimated from samples.

Closed-form $F_K(x,q)$ for Montesinos knots from Verma module recoupling in the inverted regime

Minghao Ji

Mentors: Sergei G. Gukov and Pedro Guicardi

The Gukov-Manolescu series $F_K(x,q)$ is the knot-complement avatar of the GPPV invariant: a two-variable q -series whose integer coefficients count BPS states of $T[S^3 \setminus \nu K]$. So far, closed-form expressions are available only for a few families of knots because the current method with large-color R -matrix produces state sums over infinite-dimensional Verma modules that resist resummation. We extend this significantly for all Montesinos knots. Our method replaces the R -matrix calculation by its shadow-world dual, which is graphical calculus of Verma-colored strands in which Solov'yev's quantum $3j$ and $6j$ symbols are rewritten in q -Pochhammer form. The two models compute the same invariant, but the shadow world makes fusion, and therefore the recoupling structure, obvious. This rewriting is also what allows their analytic continuation to the inverted regime.

Particle sedimentation in Newtonian and non-Newtonian fluids

Emily Jiang

Mentors: Sujit Datta, Craig Singiser, and Mohamed Abdelgawad

Sedimentation is a fundamental process governing various industrial and natural phenomena. In Newtonian fluids, particle-particle interactions generate hydrodynamic fluctuations that produce velocity fluctuations and complicate sedimentation dynamics. Viscoelastic fluids, a type of non-Newtonian fluid, exhibit both viscous and elastic properties, and further complicate settling. Sedimentation in Newtonian and viscoelastic fluids is characterized using an experimental setup and flow analysis. Glass beads are suspended in each fluid, and videos of their settling are recorded and analyzed using particle image velocimetry. Velocity fluctuations, spatial autocorrelation functions, correlation lengths, fluctuation amplitudes, and cell dependence are calculated and plotted to characterize the settling. Preliminary results indicate that viscoelastic fluids produce shorter correlation lengths and smaller velocity fluctuations than Newtonian fluids across various volume fractions. Parallel fluctuations are larger than perpendicular fluctuations, while autocorrelation trends are similar between the two fluids. These results show differences in sedimentation dynamics between Newtonian and viscoelastic fluids and establish a foundation for continuing to investigate sedimentation in non-Newtonian fluids.

From stress to the gut: Neural pathways underlying gastrointestinal dysmotility

Zhicheng Jiang

Mentors: Yuki Oka and Bochuan Teng

Stress is strongly associated with gastrointestinal (GI) dysfunction, but the neural pathways linking stress to altered gut motility remain unclear. In mice, acute restraint stress reduced upper-GI transit, increased lower-GI evacuation, and delayed feeding. Acute stress also induced c-Fos expression in the celiac-superior mesenteric ganglion complex (CG-SMG) and in stress-responsive brain regions

including the medial amygdala (MeA). Peripheral sympathetic denervation with 6-OHDA partially rescued the GI phenotype, with a stronger effect on lower-GI output, suggesting that sympathetic signaling contributes to but does not fully mediate the acute response. In contrast, broad chemogenetic activation of the MeA was not sufficient to reproduce the phenotype. Repeated restraint stress further produced a persistent reduction in baseline GI output after stress exposure ended. Together, these results suggest that stress-induced GI dysmotility arises from distributed central and peripheral mechanisms operating across acute and chronic timescales.

Localized reconstruction of electron–phonon interactions from first principles

Kehan Jiao

Mentors: Marco Bernardi and Yao Luo

This project develops a localized framework for reconstructing electron–phonon interactions from first-principles calculations. The central challenge is to retain the accuracy of a plane-wave description while obtaining a compact atomic-orbital representation that is easier to interpret and integrate into data-driven models. We first compare the reconstructed atomic-orbital band structure with the original plane-wave result to assess the fidelity of the electronic representation. We then use cross-validated learning curves to quantify how the accuracy of fitted Hamiltonian derivatives improves as additional displaced structures are included. As an independent validation of the nuclear-motion treatment, we reproduce the anharmonic phonon dispersion of palladium hydride using the stochastic self-consistent harmonic approximation. Finally, we compare the reconstructed electron–phonon coupling strengths associated with hydrogen–hydrogen, hydrogen–palladium, and palladium–palladium hopping channels and examine their temperature dependence. These results identify the local electronic processes most sensitive to atomic motion and establish a foundation for studying anharmonic and higher-order electron–phonon interactions, particularly electron–two-phonon coupling.

Quantifying the information bottleneck in radiogenomics

David L. Jin

Mentors: Lior S. Pachter and Joseph M. Rich

Radiogenomics is a field in precision medicine that integrates medical imaging data with patient genetic and molecular profiles. It provides a non-invasive way to identify tumor subtypes, predict specific genetic mutations, and choose tailored therapies. However, reported associations between imaging features and mutations in published radiogenomics papers rarely reproduce across hospitals and scanners. Our own Contrastive Language-Image Pre-training (CLIP) model failed to align the two different modalities of data and the model's predictions of the gene mutations remain random. Instead of treating every failed replication as a modeling problem, we asked a more basic question: does a computed tomography (CT) scan contain enough information about a tumor's mutations to predict them at all, no matter how good the model is?

We instead measure how much biological information survives from a somatic mutation to a CT image. Across roughly ten thousand tumors from The Cancer Genome Atlas (TCGA), mutations explain just 2% of the variation in gene expression above a calibrated permutation baseline. We then used Cohen's d to identify mutations associated with large changes in transcriptomic expression within each cancer type, yielding just 25 gene–cancer type pairs, most involving well-known tumor suppressors or oncogenes (e.g. TP53, KRAS). These genes represent the genes most likely to show up in a CT scan due to their large effect on expression. After training supervised imaging classifiers over the list of genes, we found that every single prediction came back at chance. A separate analysis using a calibrated statistical bound reached the same conclusion: CT contained very little recoverable information about the tumor's gene expression, with an estimated upper bound of 1.7% of transcriptomic variance. The strongest apparent signal was largely associated with patient sex rather than known tumor biology. Together, these results suggest that only a small amount of genetic information survives from the molecular state of a tumor to the macroscopic phenotype represented in a CT image. This provides a possible biological explanation for why many reported radiogenomic associations have been difficult to reproduce.

Developing a reusable, low-cost microfluidic device to investigate freezing fluid infiltration in porous media

Carolina L. Jones

Mentors: Xiaojing Fu and Aman Eujayl

Modeling the infiltration of freezing fluid in porous media plays an important role in diverse processes such as snow and permafrost hydrology, cryopreservation, and the design and manufacturing of advanced materials. Investigating this behavior allows for the understanding of frost heave, polar ice formation, and the mechanisms for ice nucleation and how ice growth affects the flow patterns of a liquid. Due to the natural opacity of porous media, a transparent 2D porous micromodel is needed to simplify flow imaging. To create a low-cost, reusable, and precise model, we investigated laser cutting and resin printing methods for fabrication. These methods are lower-cost and more accessible than the photolithography methods that are currently used for micromodel production. We identified a combination of 3D resin printing (stereolithography) and acrylic bonding as the most robust method to produce a precise and transparent micromodel. Using an experimental setup with a cold plate, UV lights, and a fluorescein solution, we imaged the flow in the unsaturated micromodel under subfreezing conditions. We employ binary image analysis to track the water and ice distribution in the micromodel. By systematically varying the injection rate and particle size, we aim to identify the parameters controlling the resulting flow pattern.

Optimizing FRB observations with the Deep Synoptic Array

Oliver R. Jones

Mentors: Vikram Ravi and Kaitlyn Shin

Fast Radio Bursts (FRBs) are mysterious energetic radio transients, typically with burst durations on the order of a millisecond. They are common — yet difficult to observe — astronomical phenomena, and can be utilized to investigate the energetic events that cause these transients and to probe the diffuse universe. In order to best investigate the topics above, large numbers of FRB observations and localizations are necessary for accurate population statistics. The Deep Synoptic Array (DSA) radio telescope, set to start full science operations in 2029, will provide these necessary FRB observations. In this project, we have simulated FRB populations and a DSA instrument model to understand what DSA parameters are predicted to produce the most detected FRBs. For the FRB populations, we are using the parameters acquired by the CHIME/FRB DR2 catalog to create a baseline population, and modelling how these burst parameters are affected by the DSA detection system. By altering these populations and DSA parameters (bandwidth, time resolution, frequency resolution and beam size), we should be able to optimize the number of detectable FRBs. Our simulations framework can also be used to enable further optimizing for certain sub-populations of FRBs, such as distant sources.

Investigating the role of photoreceptor-derived extracellular vesicles in early diabetic retinopathy using the TIGER mouse

William Jouayed

Mentors: Emma Lessieur and Melany L. Hunt

Diabetic retinopathy (DR) is a microvascular disease that damages the vasculature that supplies the back of the eye and is the primary cause of vision impairment and blindness in patients with long term diabetes. The vascular complications of DR have been associated with increased levels of extracellular vesicles (EVs) in vitreous samples of human donors. The cellular origin of vitreal EVs and their influence on DR are not well defined. The present study was designed to investigate if photoreceptors contribute to the pool of retina EVs using a rod-specific reporter mouse and their potential role in the development of early DR.

A 2-generation breeding scheme was followed using the transgenic inducible **GFP-EV** reporter (TIGER) mice and a rod-specific Cre recombinase (iCre75) mice. Expression of the recombinase in the progeny was screened by PCR. Photoreceptor-traceable EVs were identified by fluorescence microscopy and flow cytometry. Future work includes evaluating the influence of the transgene in the DR phenotype in a mouse model of type 1 diabetes and characterization of the molecular cargo (DNA, RNA, proteins, lipids) of GFP-tagged ret-EVs under normal and diabetic conditions.

Hormonal control of calcium isotope signals in humans: A mechanistic model of kidney handling, skeletal balance, and dietary restriction

Kyra Y. Jun

Mentors: François Tissot and Theo J. Tacail

Calcium isotopes in blood and urine have been proposed as early markers of bone loss, but recent literature shows that kidney handling and dietary calcium can dominate these signals, limiting their use as direct readouts of skeletal balance. We developed a mechanistic model in which parathyroid hormone and active vitamin D regulate calcium exchange among diet, intestine, blood, kidney, and fast- and slow-exchanging bone pools. Isotopes are propagated along these regulated fluxes, including progressive fractionation during tubular reabsorption. Under simulated dietary restriction, the model maintains blood calcium near its set point by increasing fractional intestinal absorption and renal conservation. The urinary signal shows an immediate kidney-driven response followed by a slower shift that tracks blood, suggesting that a single sample may conflate distinct physiological processes. Across simulated individuals varying in calcium intake and filtration rate, urinary isotope variability is approximately threefold greater than blood variability and decreases with increasing calcium excretion, consistent with published cohort patterns. Continuing work will further develop vitamin D storage and activation to examine deficiency, malnutrition, and skeletal growth. Paired urine and fingernail measurements will then help narrow plausible model parameters by informing short-term renal handling and longer-term isotope integration, improving inference of hormone-regulated calcium fluxes.

Long-distance time-bin entanglement distribution using dispersion compensation

Raghava K. Kalidindi

Mentors: Maria Spiropulu and Raju Valivarthi

Long distance quantum communication requires optical pulses or photons to remain distinguishable only by the information intentionally encoded in them. Chromatic dispersion in fiber broadens pulses and can reduce the interference needed for quantum networking. This project investigated dispersion compensation and Hong Ou Mandel interference using optical fibers. Coherent optical pulses were generated with intensity modulators, transmitted through fiber spools, compensated with fixed or tunable dispersion compensators, and recorded with single photon detectors and a time tagger. Pulse widths were extracted from arrival time histograms, while interference quality was evaluated from changes in coincidence counts. Classical characterization showed that pulses broadened from about 57 ps to about 96 ps after 50 km and returned to about 54 to 56 ps after compensation. Earlier interference measurements reached 47.8% visibility without fiber but decreased to 35.4% after the fiber was added. The setup was therefore rebuilt to improve frequency, timing, intensity, and polarization control. In the final configuration, the delay was held constant, and the polarization of each of the two arms was adjusted. With 50 km of fiber in the system, three recorded count comparisons produced visibilities of 47.9%, 47.3%, and 50.5%, with a total visibility of 48.8%. The recovery of an approximately 50% dip shows that polarization matching was the dominant adjustment needed in this configuration and that high visibility interference can be preserved after a long fiber spool. Repeated scans with uncertainty analysis and entangled photon sources remain necessary before the result can be used as a complete demonstration of long distance entanglement distribution.

Evaluation and analysis of the classroom experience of the Digital Kenyan Grade 10 CBC digital textbook

Jason M. Kamau

Mentors: Franca Hoffmann, Jasmine R. Bryant, and Joe Champion

INNODEMS, a Kenyan educational technology organization, has piloted an open-access digital mathematics textbook to support Kenya's newly introduced Competency-Based Curriculum. This project investigates how the textbook is used, which topics are most difficult for students, and how correctly they are actually answering questions, drawing on six months of server access logs together with structured interviews with teachers and students. Because the logs alone could not show whether an answer was correct, a difficulty proxy built from retry counts was checked directly against real scores recovered from the grading system, and found to be weak and, once contamination in the data

was accounted for, actively misleading. The structured interviews proved far more trustworthy: every topic teachers and students independently named as difficult was confirmed, or at minimum never contradicted, by this real correctness data, with reflections on geometric figures emerging as the single hardest topic in the curriculum across every method applied. This project's central conclusion is that qualitative accounts from teachers and students are a more reliable foundation for identifying difficulty than automatically logged usage data alone, and the analysis pipeline and recommendations this work produced - shared with the STACK development team and with INNODEMS - are intended to give INNODEMS an evidence-based basis for revising this textbook's content and to be reused across INNODEMS' other digital textbooks and future STACK-based deployments more broadly.

A population-level study of changing-look active galactic nuclei

Nora K. Kane

Mentor: Matthew J. Graham

Actively accreting supermassive black holes appear as active galactic nuclei (AGNs). AGNs are classified as Type 1 or Type 2 by whether their spectra show broad emission lines, a difference long attributed to viewing angle. Changing-look AGNs (CLAGNs) switch types within a few years at a fixed viewing angle. This suggests a change in accretion, possibly near a poorly constrained "few percent" critical Eddington ratio. Because CLAGNs are identified spectroscopically, yet many AGNs have only one spectrum, this work tests whether abundant photometry can reveal CLAGN behavior. Zwicky Transient Facility and Catalina Sky Survey light curves were converted to Eddington-ratio curves for 862 CLAGNs and 3,448 redshift-matched control AGNs. CLAGN and control Eddington ratios differ significantly ($p < 10^{-35}$) and CLAGNs are more variable ($p < 10^{-15}$). The threshold crossed by the most CLAGN curves places the critical ratio near 0.018, possibly decreasing with black hole mass. Ranking AGNs by recent Eddington-ratio behavior yielded one new spectroscopically confirmed CLAGN, SDSS J142125.67+394328.8. Machine-learning subgrouping separates quiescent CLAGNs, mostly below the critical ratio, from variable CLAGNs in which the relation between Fe II strength and Eddington ratio breaks down. Photometric Eddington-ratio curves can therefore flag and characterize CLAGNs, motivating spectroscopic follow-up.

Dual-space orbital learning for electronic band structure prediction of periodic materials

Minhyuk Kang

Mentors: Anima Anandkumar and Beom Seok Kang

Despite the importance of electronic band structures in computational chemistry, direct learning of full band dispersions remains underexplored, particularly at high levels of theory such as HSE06 or G0W0 quasiparticle methods. Here, we introduce OrbNet-Crystal, an orbital-informed framework predicting higher-fidelity electronic structures of crystals. Our approach employs dual-space orbital learning, combining low-cost orbital features in real and reciprocal space to capture both local chemical interactions and global band dispersion. The model predicts band structures at arbitrary levels of theory by learning corrections to a relatively computationally efficient PBE baseline.

We demonstrate the effectiveness of OrbNet-Crystal on more than 4,000 spin-resolved HSE06 band structure samples from three-dimensional crystalline materials, achieving a state-of-the-art band-energy mean absolute error of 0.18 eV. We additionally trained the model with beyond-DFT quality G0W0 quasiparticle band structures on more than 300 two-dimensional materials, achieving a strong band-energy mean absolute error of 0.09 eV. Finally, we extrapolate the model on previously unseen experimentally characterized perovskite materials testing on reported carrier effective masses and band gaps, demonstrating strong generalization beyond the training distribution. Overall, OrbNet-Crystal provides a scalable and physics-informed framework for predicting high-fidelity electronic band structures across multiple levels of electronic structure theory.

Serotonin dynamics in the VMHdm track social safety context but not valence

Zara H. Kanold-Tso

Mentors: David J. Anderson and Takumi Itakura

Serotonin (5-HT) is a neuromodulator which modulates behavioral states such as anxiety and threat evaluation. 5-HT is produced primarily by neurons in the dorsal raphe nucleus, some of which project to the dorsomedial subdivision of ventromedial hypothalamus (VMHdm). It has been demonstrated that the VMHdm is a key node in predator defense circuits wherein 5-HT tracks the current safety context, which is inherently positively valenced. However, it remains unclear whether VMHdm 5-HT dynamics specifically encode safety or instead reflect positive valence more generally. Since social interaction can be either positively or negatively valenced depending on the identity and behavior of the partner, we aimed to characterize how the serotonin-based safety state is modulated by social context. We recorded VMHdm 5-HT dynamics by fiber photometry using a genetically encoded sensor across a range of different social contexts. We also tested whether positively valenced items impact 5-HT dynamics. We found that 5-HT signal decreases upon interactions with a novel mouse, and scales with the danger of the intruder. Surprisingly, 5-HT signals decreased upon biting both positively valenced and neutral objects. These findings support the notion that 5-HT in the VMHdm represents a true safety signal rather than abstract positive valence.

How well and how soon can LISA localize EMRIs on the sky? Parameter estimation within the global fit

Thirtha Karmakar

Mentors: Katerina Chatziioannou and Michael Katz

The Laser Interferometer Space Antenna (LISA) will observe millihertz gravitational waves from Extreme Mass Ratio Inspirals (EMRIs): stellar-mass compact objects spiraling into massive black holes over 10^4 – 10^5 orbits. Because LISA data will contain many overlapping signals, all sources must be analyzed simultaneously through a “global fit.” This project scrutinizes the state-of-the-art LISA global-fit infrastructure, focusing on one specific question: are the EMRI sampling operations within the global fit working correctly? Working with a common simulated LISA data set, we constructed and validated an end-to-end Bayesian parameter estimation pipeline that combines fast Kerr eccentric waveforms (FastEMRIWaveforms), a second-generation time-delay interferometry detector response, a wavelet-domain likelihood, and parallel-tempered Markov chain Monte Carlo sampling over a 12-dimensional parameter space. In the course of this validation we identified and corrected significant bugs—most notably a sign error that had silently disabled parallel tempering—after which the sampler explores the multimodal EMRI posterior as expected. We now have visualized posterior distributions for EMRI sources analyzed at different observation times, which show sky-localization uncertainty contracting as more of the inspiral is observed, with weak, prior-dominated constraints for low-SNR sources at early times. In the remaining weeks we will produce converged posteriors for several catalog sources at multiple observation times and quantify how the sky-localization area $\Delta\Omega$ evolves over the inspiral, information that is directly relevant to prospects for electromagnetic follow-up of EMRIs.

Affective polarization across states and time: A Bayesian analysis of ANES surveys, 1964 - 2024

Shyna A. Kashi

Mentors: R. Michael Alvarez and Raquel Centeno

Affective polarization refers to the tendency to evaluate one’s own political party positively and the opposing party negatively. This project examines how affective polarization varies across states, demographic groups, and election cycles using American National Election Studies data from 1964 through 2024. I measure affective polarization as the difference between respondents’ feeling-thermometer ratings of their own and opposing parties. I address missing data through multiple imputation and estimate a Bayesian multilevel model with demographic, state, and year effects. The results indicate that affective polarization has generally increased in recent election cycles. Demographic relationships also changed substantially across decades. Older respondents consistently exhibited greater polarization than younger respondents, with the difference increasing sharply during

the 2020s. The racial pattern changed most dramatically. White respondents were substantially less polarized than non-White respondents in earlier decades, but this gap narrowed over time and reversed during the 2020s, when White respondents became more polarized. Differences by party, education, gender, and income also varied over time. These findings demonstrate that national increases in affective polarization have not occurred uniformly across demographic groups.

TIC 88297141Ab: A validated transiting planet around a young M dwarf in Upper Scorpius

Gurmeher Kathuria

Mentor: Luke Bouma

We report the discovery and validation of TIC 88297141 Ab, a transiting exoplanet around a young M dwarf member of the Scorpius-Centaurus OB Association. The host Star is 15–20 Myr old, making TIC 88297141 Ab one of the youngest validated planetary systems known. We identified the transit signal using a custom TESS search pipeline designed to detect transits in the presence of strong stellar variability, and confirmed the on-target nature of the signal with ground-based photometry (LCO/Swope and LCOGT/Sinistro). We characterized the system through a joint transit fit incorporating Gaussian processes to model stellar activity, yielding an orbital period of 4.644 days and a planet radius of $7.20 R_{\oplus}$. Speckle imaging (Gemini-S/Zorro) rules out nearby stellar companions above the contrast curve, and radial velocities from the Planet Finder Spectrograph (PFS) are consistent with stellar activity, placing an up-per mass limit of $M \sin i < 4.74 M_J$ at 2σ . Statistical validation via TRICERATOPS yields a false positive probability of $FPP = 3e-6$, successfully meeting the planet validation requirement. TIC88297141 Ab joins a small but growing sample of planets with ages below 30 Myr, offering a rare opportunity to study planets early in their lives, before significant atmospheric and orbital evolution has occurred.

Design and fabrication of in-situ loading frames for full-field deformation measurements on metamaterial structures

Daavleen Kaur

Mentors: Angkur Shaikeea and Rosan Naik

The mechanical response of metamaterials is strongly governed by their geometry. Current understanding is largely based on global structural response, but capturing micromechanical behavior requires local measurements — full-field deformations. The objective of this work is to collect 3D scans (tomographs) of metamaterials under loading (compression/tension). Since each tomograph requires a full rotation of the specimen while under load, a standard loading frame cannot be used — its columns interfere with the scan whenever they fall within the line of sight. This makes the mechanical design, not just the loading capability, a critical consideration.

Two loading rigs were designed for this program: a 1 kN rig (continuous loading) and a 2 kN rig (interrupted loading). The 1 kN rig is built to enable continuous scanning without interruption. In this mode, only one frame is captured per loading step, depending on exposure time and loading rate, so each tomograph is reconstructed from a limited number of projections. Since standard reconstruction algorithms cannot handle such sparse-projection data, a NeRF-based approach is used to generate tomographs at each loading step, followed by Digital Volume Correlation (DVC) to extract full-field deformation measurements.

The 2 kN rig uses a worm-gear drive to achieve higher loading capacity in a compact setup. It is manually operated: one handle revolution produces approximately 1 mm of plate displacement, so each loading increment corresponds directly to one full 360° scan in the XCT machine.

Biochemical reconstitution and structural characterization of the *Alvinella pompejana* inner ring complex of the nuclear pore complex

Isabella G. Kedikian

Mentors: André Hoelz and Dayeon D. Kang

Nuclear pore complexes (NPCs) are the sole gateways facilitating selective, bidirectional transport of folded proteins and ribonucleic acids across the nuclear envelope. Obstacles arise in structural studies of the NPC due to the extensive intrinsically disordered regions and high flexibility within the inner ring complex (IRC). In this project, we leverage the thermophilic annelid *Alvinella pompejana* as a model organism for IRC reconstitution because of its adaptations to increase the thermostability of its proteins. We successfully biochemically reconstituted full-length *A. pompejana* inner ring nucleoporins, CNT•Nup93, in vitro. This involved the expression and purification of full-length homologs of the CNT and Nup93 in *E. coli* and testing whether they formed defined multimeric complexes. In addition, we biochemically validated current Cryo-EM structures of the Nup205•Nup93R2 complex via 5-Alanine scanning. By mutating various residues in Nup93 with alanine, we were able to determine residue regions of mutants A1, A3, A4, and A5 play a role in this key scaffold interaction. Future steps include specific mutations of hydrophobic and electrostatic residues in these regions to observe any changes in interaction. Another focus would be the purification and reconstitution of other *A. pompejana* inner ring homologs, to achieve a complete reconstitution of the IRC.

Neural operator trajectory prediction for multi-car autonomous racing

Dami Kim

Mentors: Soon-Jo Chung and Gilchrist X. Johnson

Autonomous racing pushes vehicle control to its physical limits, and racing against other cars adds a challenge single-car methods cannot resolve: the best path for one car depends on the choices of its competitors. Reliable planning therefore requires predicting how opponents will move during close racing. This project develops and evaluates such a prediction system, built from recorded laps in a high-fidelity racing simulator. Predictions are produced in two stages. A baseline model first estimates an opponent's path from vehicle dynamics and the geometry of the ideal racing line; a learned component then corrects that estimate using the relative positions, speeds, and spacing of surrounding cars. This structure confines learning to the interactions between cars rather than the underlying motion, which improves data efficiency and keeps predictions physically plausible. Completed work includes a data processing pipeline, a set of baseline predictors, an evaluation framework based on average path error, and a real-time implementation running inside the simulator's control system. Sensitivity analysis confirms the learned component responds primarily to nearby cars rather than defaulting to single-car behavior. Remaining work is a controlled comparison against the baselines to quantify the accuracy gained during close racing.

Mod-2 negligible cohomology of groups of order 16

Matthias K. Kim

Mentor: Matthew M. Gherman

Let F be a field, G a finite group, and M an abelian group with trivial G -action. The group cohomology $H^*(G, M)$ provides information about the external structure of G . If we study the group cohomology of the absolute Galois group Γ_K of a field K the Galois cohomology $H^*(\Gamma_K, M)$ reveals information about generalized field extensions of K by M . Relative to finite group cohomology, Galois cohomology can be difficult to compute. However, since the absolute Galois group of a field is profinite, the Galois cohomology is the colimit of all group cohomology of finite quotients of Γ_K . An element $u \in H^*(G, M)$ is negligible over F if for all field extensions K/F and all Galois extensions L/K with Galois group G , u is in the kernel of the induced homomorphism $*$: $H^*(G, M) \rightarrow H^*(\Gamma_K, M)$. Hence, the more negligible classes, the "worse" our approximation of Galois cohomology. In this project, we extend the recent computation of the mod-2 negligible classes of all groups of order 8 to groups of order 16. In particular, we compute the mod-2 negligible classes of all abelian groups of order 16 over formally real fields and fields containing sufficient roots of unity. We plan to study cohomological invariants to obtain results for the remaining abelian and non-abelian cases over a general field.

Development of a biomimetic intervertebral disc implant using digital light processing 3D printing and wet-spinning

Myu Kim

Mentors: Julia R. Greer, Caroline Dial, Jingtong Zhao, and Julia A. Kornfield

Intervertebral discs (IVDs) are fibrocartilaginous joints located between adjacent vertebral bodies that provide mechanical support and spinal mobility. Each IVD consists of the nucleus pulposus, annulus fibrosus, and cartilaginous endplates. IVD degeneration is recognized as a leading cause of lower back pain. Current clinical treatments are largely limited to pain management or invasive surgical interventions, such as spinal fusion, which permanently alter spinal mechanics rather than restoring natural joint function. Furthermore, existing synthetic implants fail to replicate the complex architecture and biomechanical properties of a healthy disc. To address these limitations, we established a biomimetic approach to fabricate whole-disc implants. Utilizing methacrylated fibrinogen (FibMA) blended with polyethylene glycol diacrylate (PEGDA) and collagen fibers, we developed a fabrication strategy that combines digital light processing (DLP) and collagen fiber wet-spinning. This multimaterial fabrication approach enables control over scaffold architecture and mechanical properties to reproduce the distinct structural and functional characteristics of native IVD components within an assembled construct. Preliminary data indicate that these materials support cellular differentiation toward IVD-specific phenotypes. Together, this approach produces biomimetic whole-disc implants whose component-specific mechanics guide cell function and promote a native tissue phenotype.

Characterizing thermohaline convection in polluted white dwarfs using 3D hydrodynamical simulations

Marcus N. King

Mentors: James W. Fuller and Jose Rafael Fuentes

Recent surveys have revealed a large white dwarf population with evidence of polluting exoplanetary material in their atmospheres. Given stellar conditions, these materials should have very short gravitational sedimentation timescales, indicating ongoing accretion from debris disks onto the stellar envelope. In the envelope, this drives thermohaline (fingering) convection. Improved characterization of these materials requires a quantitative understanding of this mixing process. Recent work has established analytic solutions for thermohaline-unstable convection zone propagation within stellar interiors. We computationally verify these solutions via a series of 3D Boussinesq hydrodynamic simulations with varying initial conditions and fluid parameters, using the Dedalus code. We focus on the dynamics of this physical system in unmixed and premixed simulations. We find both model dynamics supporting analytical assumptions within previous work, and we also see a strong coupling between $h(t)$ and $X_s(t)$ in each case, supported by calculated power laws. These results demonstrate that the time-dependent propagation of the thermohaline-unstable front, coupled with surface heavy element abundance, may provide insights into the accretion history of these polluted white dwarfs.

Using reaction network structure to recover metabolic rates from carbon labeling measurements

Elianna G. Kondylis

Mentors: Gregory Stephanopoulos and Rustem F. Ismagilov

Cells run chemical reactions at rates that cannot be measured directly. These rates are instead inferred by feeding cells a carbon-13 labeled nutrient and measuring how the label spreads through downstream molecules. Established software recovers rates from these patterns by repeated testing, which is slow and can settle on unstable solutions. A recent neural network method is far faster but treats each measurement as an independent number, discarding all knowledge of which reactions connect which molecules. This project asks whether restoring that knowledge improves the recovery of reaction rates. The labeling network is represented as a hypergraph in which one connection may join three or more molecular fragments at once. This form is necessary because some reactions fuse two separate molecules into a single product, so the product's labeling depends jointly on both sources and cannot be expressed by connections that link only pairs. Fragments serve as nodes and carry measured labeling patterns; connections carry the unknown reaction rates. Repeated rounds of

averaging along the connections let measured fragments inform unmeasured ones. This model can help with rate predictions using stationary and, in a novel addition, for isotopically non-stationary data.

Engineering cpSRP43 for enhanced chaperone activity against amyloid aggregation through directed evolution

Fanhao Kong

Mentors: Shu-ou Shan and Bryce Yeazell

Chloroplast signal recognition particle 43 (cpSRP43) is evolved to enhance its chaperone activity toward amyloidogenic aggregation-prone proteins, including α -synuclein and amyloid- β (A β), associated with neurodegenerative disease. The L189R mutation—a key allosteric residue predicted to stabilize the closed conformation of cpSRP43—was introduced into four previously evolved tailor-made molecular chaperones (TMMCs), and circular dichroism showed that L189R destabilizes the closed conformation of cpSRP43 at physiological temperature, and seeded α -synuclein fibrillation assays showed that L189R broadly impairs anti-fibrillation activity. In parallel, a single-region saturation-mutagenesis library targeting seven substrate-binding regions of cpSRP43 was constructed and is being screened against A β 42 by HigB/HigA Toxin-Antitoxin Chaperone Selection (HITACS). Variants beginning to emerge from this screen will inform combinatorial library design to probe epistasis and will serve as starting points for continuous directed evolution via the Orthogonal Replication System (OrthoRep).

A reduced-order thermal model and multi-input controller for the HISPEC spectrograph

Andrey Korolev

Mentor: Jake Zimmer

Cryogenic astronomical instruments such as the HISPEC spectrograph must hold large optical assemblies at a stable operating temperature near 55 K, yet their thermal behavior is set by hundreds of thousands of coupled conduction and radiation pathways, far too many states to simulate in real time or to design a controller against. This project develops an end-to-end pipeline that turns a CAD assembly into a lumped thermal-network model and reduces it to a compact form suitable for on-board control. From the instrument geometry an octree mesh builds a multi-million-node thermal graph; the linearized dynamics are projected onto their 120 slowest physical modes and then reduced by balanced truncation to a 30-state model that keeps only the jointly controllable and observable dynamics. On this reduced model I design a multi-input linear-quadratic-regulator controller with a static state estimator and a regularized steady-state feedforward for its 33 heaters and 28 controlled sensors. The 30-state model reproduces the full plant's heater-to-sensor step response to within about 0.15%, small enough to design against while remaining light enough to run on a microcontroller. Remaining work adds surface-to-surface radiation and hardware deployment.

Investigating how gene-environment interaction causes heart defects

Aleen A. Koujakian

Mentors: Duncan Sparrow, Marianne Bronner, and Victoria Rashbrook

Congenital heart disease (CHD) is a non-communicable disease in which the embryonic heart fails to form correctly. CHD is the most common class of human birth defects, but only one-third of cases are attributed to genetic causes. Maternal iron deficiency anemia (IDA) has been identified as an environmental factor that causes CHD in mice. *NOTCH1* haploinsufficiency is also associated with CHD but has high phenotypic variability, which may be explained by gene-environment interaction (GxE) between the *NOTCH1* null allele and maternal IDA. This GxE has been demonstrated experimentally in mouse embryos heterozygous for a global *Notch1* null allele by an unknown mechanism. Here, we investigated the molecular basis of this GxE using whole-mount RNAScope and immunofluorescence to detect changes in expression of genes and proteins required for embryonic heart development. We also sought to identify the cell populations that require *Notch1* expression by phenotyping mouse

embryos with conditional *Notch1* deletions in the anterior second heart field, neural crest, or endothelium. High resolution episcopic microscopy (HREM) was used to create 3D models for manual and automated morphology analysis. This work may elucidate how maternal IDA interacts with embryonic genetics to cause CHD and inform future studies aimed at lowering CHD risk.

Toward a mechanistic understanding of Daple regulation by cytoplasmic dynein-1

Aurelia H. Kuester

Mentors: Aga Kendrick, Isaac Fianu, and Alvaro de la Gandara

The spatiotemporal regulation of signaling pathways via G proteins is essential for healthy cellular function. In addition to G protein-coupled receptors, non-receptor activators have emerged as important regulators of G protein signaling. One such activator, Daple, facilitates guanine exchange in G proteins, thereby regulating noncanonical Wnt signaling, a pathway critical for cell polarity, migration, and cancer progression. However, the mechanisms governing Daple's regulation, subcellular localization and signaling activity remain poorly understood. Recent work from the Kendrick lab identified Daple as a cytoplasmic dynein-1 (dynein) activating adaptor, potentially linking dynein transport to cell signaling. This project investigates the structural basis of Daple-G protein binding as well as the functional consequences of Daple-dynein interactions on G protein activation. Using bioluminescence resonance energy transfer (BRET) assays and cross-linking approaches we examined how mutations that disrupt Daple-dynein interactions affect the separation of G protein subunits. We found that disruption of the Daple-dynein interaction reduced Daple-mediated G protein activation, while optimization of BRET conditions enabled consistent detection of this activity. Additionally, we established conditions for cross-linking Daple-Gai3 complexes, providing a foundation for future structural analysis of this interaction. These advances provide an experimental pipeline to define how dynein-mediated transport spatially regulates Daple-mediated G protein activation.

Classification of alerts from the Roman Space Telescope

Aayush Kuloor

Mentors: Ashish Mahabal and Jacob Jencson

The Nancy Grace Roman Space Telescope will produce an estimated 12.5 billion time-domain alerts, making manual inspection impossible. This project develops machine learning classifiers for the Roman Alerts Promptly from Image Differencing (RAPID) Project Infrastructure Team, across survey environments with varying depth, crowding, filters, and observing cadence. The goal is a unified framework that first rejects bogus detections and then classifies genuine astrophysical sources. Using simulated Roman survey data from the High Latitude Time Domain Survey (HLTDS) and Galactic Bulge Time Domain Survey (GBTDS), the models combine image cutouts, point source function (PSF) measurements, survey and filter information, and forced photometry light-curves. The artifact rejection model uses a rotation-invariant convolutional image encoder, with cross-attention between image features and tabular metadata. The source classification model uses a transformer-based light-curve encoder with continuous time embeddings to separate transients from variable sources. Evaluation used held-out detections and object separated splits to reduce source leakage between training and testing. The final framework includes bogus filtering, transient/variable classification, positive-unlabeled learning for uncertain labels, and controlled ablations measuring the value of light-curve information. The results show that image and PSF features are effective for artifact rejection, while temporal light-curve information substantially improves source classification.

Growth and transport characterization of thin β -FeSe via confined flux synthesis

Vijay R. Kumar

Mentors: Linda Ye and Gabriel Perko-Engel

The realization of fault-tolerant quantum computing relies heavily on materials capable of hosting topologically protected quantum states. Currently, researchers are particularly interested in iron-based superconductors. Specifically, thin β -FeSe has emerged as an ideal platform due to its highly tunable critical temperatures (T_c) and its potential to host Majorana zero modes. However, current growth techniques (e.g. MBE, CVD, ALD, etc.) fail to fabricate high quality thin crystals with smooth surfaces, large area, and low stoichiometric impurities. We employ a NaCl/KCl bulk flux technique confined

between mica layers to fabricate large area, high quality thin films of β -FeSe. Using structural and electronic characterization (e.g. Raman spectroscopy, transport, X-Ray Diffraction, and Atomic Force Microscope) we aim to confirm surface uniformity, elemental composition, and high temperature superconductivity.

GPU-accelerated multicoil Bloch simulations for nonlinear inversion in MRI reconstruction and quantitative mapping

Alya Kurt

Mentors: Shreyas S. Vasanawala, Andres Collazo, Carlos A. Castillo Passi, and Daniel Ennis

Accelerated MRI requires reconstruction methods that use receive-coil sensitivities. Model-based reconstruction methods have shown improved image quality and accuracy by incorporating physical effects, but need sequence-specific adjoints. Differentiable physics-based reconstruction maintains image quality while being reusable across sequences; but at a high computational cost. To develop practical differentiable physics-based image reconstruction we: (1) developed a GPU-accelerated simulation which includes multi-coil sensitivities by extending KomaMRI.jl, and also (2) used this simulation in an optimization problem to retrieve parametric maps like spin density, T1 and T2, where Enzyme.jl and Reactant.jl were used for gradient calculation via reverse-mode automatic differentiation (AD). In this way, we can ensure data consistency, as the predicted KomaMRI-simulated signals are forced to match target measurements. We compared our approach using AD against finite differences (FD). FD took 60 hours to reconstruct in comparison with our AD-based framework that completed the same 12-coil nonlinear-inversion in 7 minutes. Furthermore, compared to traditional reconstruction methods like CG-SENSE, we achieved an error of 0.422% (normalized data-consistency residual) vs 6.05% for CG-SENSE, representing a 93.0% error reduction. Thus, this work enables fast and accurate physics-based MRI reconstruction in a pulse-sequence agnostic way.

Twisted CSS codes

Colin H. La

Mentors: Andrew Lucas and Xie Chen

Calderbank–Shor–Steane (CSS) codes form the basis of many quantum low-density parity-check (qLDPC) code constructions. Despite significant progress in constructing qLDPC codes with favorable rate and distance scaling, fault-tolerant logical gate implementation has remained a challenge. In particular, several no-go theorems impose severe limitations on transversal gates, especially in the presence of geometric structure. Twisted CSS codes generalize ordinary CSS codes by decorating X-checks with diagonal phases, circumventing some of these logical gate implementation constraints. We construct a general algebraic framework for twisted CSS codes which characterizes twisted checks, the constraints imposed by their commutation relations, the codespace, and logical operators. Additionally, we propose an efficient two-stage decoder that first corrects X-errors and then addresses remaining diagonal errors, thereby mitigating the twist-dependent errors induced by measuring twisted X-checks in the presence of residual X-errors.

Towards precise simulation of molecules with memory-efficient analytical gradients from auxiliary-field quantum Monte Carlo

Daniel S. Lai

Mentor: Sandeep Sharma

Molecules can be simulated to near chemical precision using electronic structure theory in quantum mechanics. For example, computing the gradients of the total molecular energy with respect to certain variables enables the prediction of many useful properties, such as the geometric structure and dipole moments of a molecule. Auxiliary-field quantum Monte Carlo (AFQMC) stands out as a powerful and scalable approach that is a rising alternative to traditional quantum chemistry methods. The current approach to reliably compute energy gradients from AFQMC relies on standard reverse-mode automatic differentiation, which is analytical and time-efficient but incurs a significant memory overhead. We explore a new analytical strategy for energy gradients with low memory cost comparable to the evaluation of the energy itself, which could extend the current capability of high-precision simulations of larger molecules.

Breaking down the concentration-dependent dielectric response of lithium battery electrolytes using atomistic simulations

Chung Wang Lam

Mentors: Kara D. Fong and Nidhish Lella

Lithium-ion batteries find themselves in various areas of application, from small personal electronic gadgets to mass energy storage solutions and aerospace engineering. The dielectric constant, derived from the general dielectric response of the medium, is a crucial thermodynamic parameter of battery electrolytes, as it controls the strength of all electrostatic interactions. In turn, it affects the effectiveness of ionic conduction in the electrolyte. In accurately predicting properties of battery electrolytes, the dependence of dielectric constant on the battery salt concentration needs to be understood. So far, experimental measurements through dielectric spectroscopy have confirmed a non-monotonic dependence, but its physical origin is still unclear. Here, we utilise molecular dynamics simulations to study such a phenomenon in the cases of LiAsF_6 in DMC and LiPF_6 in 3:7 (by weight) EC:EMC. By studying the separate contributions to the dielectric response by solvent-solvent, ion-solvent, and ion-ion interactions, we provide an atomistic understanding of the physical mechanisms underlying the concentration-dependent dielectric response in terms of the competition between two effects: the formation of ion pairs and the destruction of solvent dipole correlations by free ions. The result provides insight into the design and engineering of similar lithium-based electrolytes.

What to reward and where to start: Guiding robots that learn by trial and error

Jadon T. Lam

Mentors: Pietro Perona and Ilona Demler

Robots can learn skills in simulation by trial and error, guided by a numerical score, the reward, that tells them how well they did. When the reward is given only for success and success is rare, the learner receives almost no guidance. Two common remedies are to reshape the reward so that it also pays for progress, or to start practice attempts partway through the task, at states taken from a demonstration. This project examined both. In a simulated block-stacking task, a standard shaped reward trained a seven-joint robot arm to bring a cube to its target position in about 90% of attempts. In this first training run, however, a stricter measure requiring the cube to be released and at rest recorded no completed stacks. The reward multiplies progress by the gripper's closeness to the cube, so letting go forfeits most of it. In a second, controlled study, a simulated rocket had to cross a maze with four narrow gaps and was rewarded only for reaching the goal in time. Starting practice from 128 points along a single demonstration made the task learnable: 16 of 16 training runs succeeded, compared with 0 of 5 without these starting points. Unlocking the points backward from the goal made learning about twice as fast. In pre-registered comparisons with up to 32 independent runs per condition, no smarter placement or learned weighting of the points did significantly better than spacing them evenly, and advantages seen in smaller experiments disappeared as the samples grew. A reward for progress along the same demonstration produced no successes in 34 runs. The results suggest that how a demonstration is used matters more than which of its states are chosen, and that shaped rewards can fail in ways a headline success rate does not reveal.

A GPU-native complex auxiliary-space Maxwell solver and high-order finite-element solver for scalable 3D magnetotelluric forward modeling

Henry B. Lane

Mentors: Paul D. Asimow and Carl Swindle

3D magnetotelluric (3DMT) forward modeling has consistently posed itself as a computationally demanding problem. Despite its wide range of applications, ranging from geothermal and mineral exploration to studies of crustal and mantle structure, the high computational cost of such simulations remain prohibitive for developing large-scale inversion methods and uncertainty quantification methods. Building on recent advances for solvers that combine high-order Nédélec-type finite elements with the auxiliary-space Maxwell (AMS) preconditioner, this work develops a complex-valued, GPU-native AMS solver and constructs a modified high-order 3DMT solver around it, eliminating the need for the real-valued block decomposition. Moreover, we validate the solver by comparing it against other validated solvers and accepted numerical benchmarks, demonstrating

accurate MT responses. Finally, we demonstrate that the solver has improved convergence behavior relative to the reference solver baseline, when implemented on GPUs inherits the favorable scaling behaviors of the CPU-based reference solver, and show that the associated wall-time is significantly reduced relative to the CPU implementation of the reference solver. Together, these developments provide a high-accuracy, geometry (bathymetry/topography) conforming, and computationally convenient and scalable foundation for next-generation 3DMT inversion.

Phenotypic screening and amplicon-seq genotyping of multiplex CRISPR mutants affecting *Arabidopsis* root development

Erik R. Ledvina

Mentors: Trevor M. Nolan and Xiaohui Li

The transition from proliferation to differentiation is essential to plant root development and thus vital to the rest of the plant's growth. A systematic understanding of regulatory factors involved in this transition will enable scientists to bioengineer more viable crops. However, because of gene redundancy in plants, studying plants via traditional methods of single loci genetic perturbation is ineffective. Thus, in this project, plant lines mutated using a multiplex CRISPR gRNA library that targets multiple genes within a gene family, were used to study fully perturbed gene function. Phenotypic screening was used to identify promising mutants that displayed short root length and next-generation sequencing was used to sequence gRNA containing amplicons to identify the multiplexed target genes. Finally genotyping the target loci of potential novel mutant candidates will be used to support that mutant phenotypes were not the result of off-target effects. Through this forward genetic pipeline, sequencing has revealed mutants with potentially novel or understudied genetic mutations. Future work will pertain to interpreting sequencing results, finishing genotyping, and ordering T-DNA insertion mutants to further support the mutant phenotype.

Development of a coarse-grained force field for perfluorononanoic acid (PFNA) micellization: From quantum mechanics and molecular dynamics to coarse-graining

Hsin-tien Lee

Mentors: William A. Goddard III and Sejun Kim

Per- and polyfluoroalkyl substances (PFAS), such as perfluorononanoic acid (PFNA), are highly persistent environmental pollutants that pose serious risks to public health and ecosystems. In aqueous environments, PFNA molecules self-assemble into micelles, a process that influences remediation strategies. A major challenge in computationally studying these systems is that standard force fields struggle to accurately describe the highly electronegative perfluorinated backbone. In this work, we aim to develop a multiscale computational model, transitioning from quantum mechanics (QM) and molecular dynamics (MD) to a coarse-grained (CG) force field, to investigate PFNA micellization. Through a systematic QM analysis of the torsional energy landscape, our ongoing efforts focus on modifying atomistic force fields to better describe the molecule. Ultimately, we aim to develop a coarse-grained model for PFNA, enabling large-scale simulations of PFAS self-assembly and phase behavior.

Chromatin accessibility and cofactor rules predict Runt regulatory behavior of the upstream *net* enhancer

Jayden D. Lee

Mentors: Angelike Stathopoulos and Isaryhia M. Rodriguez

Runt is a transcription factor most studied for its role as a segmentation gene critical to *Drosophila melanogaster* development. A few well-studied segmentation genes establish Runt as both a repressor and activator, and indicate this is influenced in part by cofactor presence. However, little is known about Runt regulation across the genome outside of these genes, as well as what rules govern Runt's context-dependent regulatory behavior. Thus, we are interested in whether genomic data and cofactor context can predict Runt site-specific behavior. To answer these questions, we leveraged ATACseq data showing chromatin accessibility differences between Runt mutant and wildtype embryos as well as ChIPseq data showing Runt DNA occupancy to support prediction of genes likely to be directly bound and potentially regulated by Runt. We then used *in situ* hybridization of enhancer-reporter

transgenic lines to functionally validate our predictions. Finally, we assessed the genome-wide extent of Runt regulatory behavior and how it is affected by known cofactors. Here we show that genomic data predict Runt activating behavior at the upstream enhancer of *net*. These results support that chromatin accessibility and cofactor context may provide predictive signatures of Runt regulatory behavior at previously uncharacterized loci.

Electronics design and sensor calibration of the magnetosonic transient plasma analyzer for characterizing hypervelocity impact plasmas on a lunar mission

Jeylin E. Lee

Mentors: Sigrid Elschot, Glen A. George, and Neilabjo Maitra

Hypervelocity impacts of meteoroids with spacecraft surfaces may generate dense plasma clouds, which then expand, propagating radio frequency (RF) emissions and instabilities that are hypothesized to cause catastrophic system anomalies. The magnetosonic transient plasma analyzer (MTPA) is being developed to measure the particle energy and velocity distribution functions of different ion species in order to understand the underlying plasma physics and evaluate the mechanisms for harmful RF emissions. The MTPA will eventually fly on a satellite or lunar orbiter to characterize the plasmas generated in the space environment, thus there is a need to adapt current prototypes into flight-compatible electronics. Data acquisition and power distribution electronics have been redesigned and miniaturized to fit eight signal channels on a single PCB, rather than eight separate boards. Firmware upgrades were written in C, and can be flashed onto the RP2040 microcontrollers over the new SWD interface. Experiments using legacy plasma sensors in the laboratory vacuum chamber were performed to begin calibration of the MTPA, using a high-voltage source to induce dielectric breakdown and generate plasmas. Next steps involve assembly and integration into a functional device that can be easily deployed for field tests, such as at NASA Ames Vertical Gun Range.

Experimental validation of distributed-aperture beam synthesis and characterization of the SKIPR telescope for superconducting millimeter-wave imaging

Maya J. Lee

Mentor: Jack Sayers

Millimeter-wave imaging systems operating at 90-150 GHz provide long-range imaging capabilities through atmospheric obscurants but remain fundamentally limited in angular resolution by aperture size. This project investigates distributed-aperture beam synthesis while simultaneously characterizing the optical performance of the Superconducting Kinetic Inductance Passive Radiometer (SKIPR) millimeter-wave telescope. Laboratory measurements of a dual-aperture optical system successfully reproduced previously demonstrated beam patterns, validating the experimental setup through comparison with electromagnetic simulations and preparing the system for future telescope integration. In parallel, telescope characterization efforts have continued using beam maps collected at source distances of approximately 100 m and 3 km. Following improvements to the telescope optics that reduced stray-light contamination during beam mapping, current work focuses on understanding the telescope's operational behavior across multiple source distances and identifying the transition from near-field to far-field operation. Optical simulations in Zemax OpticStudio and electromagnetic simulations in Altair FEKO compare predicted and measured beam profiles and evaluate the telescope's operational limits. Ongoing software development focuses on detector position calibration, source isolation, and beam-map analysis to enable direct comparison between measured data and numerical simulations. Together, these efforts improve telescope operation, support distributed-aperture observations, and guide calibration strategies for long-range superconducting millimeter-wave imaging systems.

Spectral analysis of supersonic cylinder wake

Taehan Lee

Mentors: Joseph E. Shepherd and William Feasey

Supersonic cylinder wake is a foundational case of compressible bluff-body flow. To image and analyze the wake flowfield, this project utilizes high speed schlieren imaging in combination with Spectral Proper Orthogonal Decomposition (SPOD) to characterize the acoustic mechanisms within the flow. Data was collected in the Caltech Ludwig Tube, a blowdown Mach 4 facility particularly suitable for spectral analysis due to its extended test-time capabilities. An extended cylinder geometry mount was designed to improve on previous work. By varying the upstream driver pressure and cylinder diameter, we match a series of Reynolds numbers ranging from $8.9E4$ to $4.89E5$. Downstream slipline oscillations comprising the characteristic collapse frequency of the flow are analyzed. Additionally, upstream acoustic structures are measured supporting the existence of an aeroacoustic feedback loop.

Inverse kinematics derivation and mechanism development for thrust vector control on a self-lander

Zoe Lee

Mentors: Mory Gharib and Jack N. Caldwell

In recent years, there has been a growing emphasis on reusable, self-landing flight vehicles in the space industry, necessitating highly demanding flight profiles during the various stages of a flight, including reentry and landing. Careful maneuvers during the descent require dynamic adjustment of the thrust against significant aerodynamic and inertial disturbances. Consequently, expanding the performance envelope of robust thrust vector control (TVC) systems, both in mechanical design and control architecture, has become fundamental for enabling optimal, reliable, and predictive launch and landing operations, not only for large commercial vehicles but also for novel, smaller scale artifacts for academic research. The redesign, manufacturing, and validation of an improved TVC system for a collegiate-level self lander under development has been explored. The newly developed design underwent critical design reviews with updated inverse kinematics validation, detailed CAD development and component sourcing and machining in order to meet and exceed requirements of a robust and reliable gimbaling system. All components were subsequently assembled into the full TVC unit, subject to integrated validation progress alongside the remaining subsystems of the lander in the future.

Covariance computation of CMB lensing power spectrum

Chi Suen Leung

Mentors: Wai Ling Wu and Nicolas C. Ferree

The cosmic microwave background (CMB) provides a powerful probe of both the physics of the early Universe and the growth of large-scale structure. In particular, weak lensing by intervening matter remaps the primordial CMB temperature anisotropies and polarization, inducing statistical mode coupling that can be used to reconstruct the lensing field.

We have modeled the covariance matrix of CMB lensing power spectra reconstructed from SPT-3G, Planck, and ACT DR6 analytically. Cross-covariance between experiments is driven not only by the shared underlying gravitational potential, but also by correlated reconstruction noise. This noise correlation stems from the identical CMB fluctuations in their overlapping footprints. Our results for SPT-3G winter field show a maximum 14% total correlation with Planck and 21% with ACT. We have extended the formalism for SPT-3G summer and wide field and found a maximum correlation between SPT wide field and Planck of 26%.

Using a Gaussian likelihood built with the analytical bandpower covariance matrix, we evaluate the impact of inter-experiment correlations. Preliminary Markov Chain Monte Carlo (MCMC) results show that ignoring correlations between the SPT main field and other experiments artificially tightens the 1-sigma confidence level on Ω_m by 4%. Furthermore, in the full joint analysis of SPT Ext-10K, Planck, and ACT, neglecting these cross-correlations artificially narrows the constraint by 9%.

Mathematical models of flexible neuronal participation

Christopher H. Leung

Mentor: Carlos Lois

The mechanisms by which the brain maintains a stable output as the upstream neuronal representation changes remain incompletely understood. The HVC nucleus of the zebra finch song motor circuit demonstrates an ability to restore the original birdsong after lesioning, even though neurons participate stochastically in song generation. We jointly model the song's acoustic features with neuronal calcium imaging on the song syllable scale. Using a hierarchical Bayesian model, we confirm the role of song syntax in neuronal participation. We then statistically decode the syllables to quantify the stability of the neuronal representation even as individual neurons' participations shift after lesion. Finally, we adopt an Ising model to distinguish between network redundancy and rewiring. Our data suggest multiple solutions to the reliability-flexibility tradeoff in motor circuits: the unreliable participation of neurons may provide a flexible manifold for both learning and recovery

Neural operator splitting for Monte Carlo fluid simulation

Jianxiang Li

Mentors: Anima Anandkumar, Hong Chul Nam, and Xi Deng

Accurate fluid simulation around complex boundaries is computationally expensive, particularly for methods that use random sampling rather than geometry-conforming meshes. This project develops learned approximations of a velocity-based Monte Carlo fluid solver and investigates how the organization of learned updates affects error accumulation over long simulations. We introduce three complementary forms of neural operator splitting. The first learns separate updates for the solver's advection, diffusion, and pressure-correction stages. The second alternates models trained for successive positions in a prediction sequence, enabling later models to account for errors produced earlier. The third applies rotated or reflected versions of a learned update and then restores the result to its original orientation. The models are trained on solver-generated velocity fields and evaluated on flows with and without solid obstacles, including previously unseen airfoil shapes. Experiments indicate that all three strategies improve long-horizon prediction accuracy over repeatedly applying a single model, with their combination providing the strongest overall performance. Ongoing work extends the framework to three-dimensional and broader flow settings. Neural operator splitting thus offers a practical path toward more accurate long-range fluid prediction while preserving useful structure from the original solver.

Biomechanically driven forward neural encoding: Synthetic EEG generation

Joan Li

Mentors: Yanan Sui and Glen A. George

Electroencephalography (EEG)-based motor execution decoding holds immense promise for closed-loop assistive devices and physical neurorehabilitation. However, training high-capacity decoding architectures remains fundamentally bottlenecked by the high cost, time intensity, and physiological variability of collecting large-scale human EEG datasets. In this work, we propose a physics-grounded data augmentation framework that utilizes biomechanical simulation for neural encoding by using the MS-Human-700 model—a high-fidelity whole-body musculoskeletal simulation of the human body comprising 700 muscle-tendon units. We extract high-dimensional, continuous muscle activation forces and kinematic trajectories during simulated motor execution tasks. These physiological excitation signals are processed through a synthetic forward neural encoder to generate multi-channel, scalp-projected synthetic EEG time-series. By mapping deep musculoskeletal dynamics directly to cortical representations, our pipeline provides a physics-consistent synthetic data generator for neural decoding models. This offers a scalable solution to neural data scarcity in assistive robotics.

Toward absolute protein quantification in the mammalian circadian clock: Evaluating flow cytometry and Western blotting

Qianqian Li

Mentors: Michael B. Elowitz and Gal Manella

The circadian clock is controlled by interacting proteins and genes. Quantitative models of how perturbations alter clock behavior require measurements of relative and absolute protein abundance. This project aimed to measure ectopic protein abundance relative to endogenous protein abundance, while developing a method for absolute quantification. A single-cell flow cytometry assay was evaluated using a fused fluorescent tag to measure ectopic protein and intracellular antibody staining to estimate the total target protein pool. The two signals were compared across induction or depletion conditions. Although the fused tag clearly reflected induction, the antibody signal showed no corresponding increase. Changes to permeabilization and antibody concentrations did not change the result. Western blotting was therefore explored as a population-averaged alternative and as an independent test of antibody specificity. Preliminary experiments revealed variable antibody specificity and identified insufficient loading, low population-averaged target abundance, and suboptimal transfer as barriers to quantitative interpretation. The current flow cytometry assay does not yet support quantitative comparison. Future work will test whether the fused tag alters epitope accessibility using new cell lines and optimize Western blotting for absolute quantification with purified standards.

Adaptation of sympathetic nervous system in pregnancy

Reina E. Li

Mentors: Yuki Oka and Wongyo Jung

The sympathetic nervous system is a part of the body's autonomous nervous system and plays a critical role in maintaining homeostasis, including the regulation of blood pressure, heart rate, and blood glucose. Pregnancy establishes a new baseline for these physiological parameters by recalibrating the body-brain axis. In this project, we compare sympathetic nerve innervation by performing immunohistochemistry staining on the kidney, pancreas, and duodenum of virgin, pregnant, and lactating mice. Using ImageJ, we obtain area fractions between virgin, pregnant, and lactating kidneys and observe a significant difference between virgin and pregnant area fractions and no significant differences elsewhere. These results are interesting because it implies that sympathetic nervous system activity is elevated specifically during pregnancy, and has interesting implications for diseases that occur during pregnancy such as gestational diabetes. However, our data is still limited and more is needed before fully coming to a conclusion.

Advancing monolithic integration of vacuum tubes on CMOS

Richard Y. Li

Mentors: Ali Hajimiri and David A. Baum

Although solid-state transistors dominate modern electronics, vacuum tube devices still show superior power density and high frequency performance. Monolithically integrating vacuum tubes onto standard complementary metal-oxide-semiconductor (CMOS) substrates allows the benefits of vacuum tubes and transistors to be combined to increase the performance of radio frequency integrated circuits. This project addresses numerous constraints of current on-chip vacuum tubes, namely their short lifespans and limited gain. By optimizing the geometry of the vacuum tubes, both the lifespan and gain of the vacuum tube devices may be improved; thus, a variety of structures were designed to determine the optimal geometry for vacuum tubes. The measurement of these structures will also be used to improve models of vacuum tubes. Additionally, supporting digital control circuitry was developed to precisely control the temperature of the vacuum tubes to increase lifespan. These developments will further the goal of building a fully integrated vacuum tube system on standard CMOS chips.

How do rattlesnake rattles make sound?

Haitong Lian

Mentors: Chiara Daraio and Xiaoxiao Xiong

Rattlesnakes produce a characteristic broadband sound using their rattle, a chain of hollow, interlocking keratin shells. Although shake rate is known to modulate the acoustic signal behaviorally, how segment geometry shapes the underlying mechanical response remains poorly understood. We investigate how structural and actuation parameters influence the collision and vibrational dynamics of the chain by decomposing the problem into several stages. The shell geometry is parametrized, assembled into an interlocking chain, and simulated using the Level Set Discrete Element Method (LS-DEM), with the chain driven periodically at one end. These simulations provide the location, timing, and impulse of each collision. The collisions are then applied as impact loads to an individual shell in a finite-element model using COMSOL Structural Mechanics. The resulting vibration is well approximated through modal superposition, expressed as a sum over the impacts experienced by the shell and its vibrational modes. Finally, we use COMSOL Acoustics simulations to identify the leading modal contributions to the rattle sound and compare the simulated spectrum with the biological sound spectrum. By integrating tail-scale collision dynamics, shell-scale vibration, and acoustic propagation, this work provides a framework for understanding and designing bio-inspired interlocking structures with tunable vibrational and acoustic properties.

Guiding denoising diffusion models for structure-based drug design via machine-learned interatomic potentials

Yiyan Liao

Mentors: Anima Anandkumar, Amin Tavakoli, and Yu Zhang

Diffusion models are the dominant approach to structure-based drug design, yet trained only to fit reference structures, they often place ligands in physically implausible poses that clash with the protein. We introduce MLIPDiff, an inference-time guidance that can be applied to any pretrained diffusion models for ligand generation. The guidance is physics-informed using UMA (Universal Models for Atoms), a machine-learned interatomic potential (MLIP). The guidance is the weighted sum of four energy terms (protein-ligand interaction, ligand strain, steric clash, and pocket anchoring) and runs on the same checkpoint with no retraining. On CrossDocked2020, MLIPDiff completely eliminates protein-ligand collisions while preserving the ligand's local geometry, improving its bond-length distributions and physical plausibility over the base model. Compared to models such as NucleusDiff that are specifically designed to reduce atomic clashes, MLIPDiff reaches the same clash-free structures without conformational collapse. It maintains the PoseBusters rate compared to the base model and attains the best median Vina Score, Vina Min, and Vina Dock, surpassing both the base model and NucleusDiff, with 59.6% of ligands binding more strongly than the reference. MLIPDiff thus reaches a niche prior methods left open, physically plausible and high in affinity at once, from a single frozen model at no training cost.

Effects of kinematic reference quality on dynamics-aware humanoid motion retargeting

Yit Xueng Lim

Mentors: Aaron D. Ames, Junheng Li, and Sergio A. Esteban

Creating kinematically feasible robot motions is an important first step in humanoid motion-generation pipelines, but the extent to which reference quality affects subsequent dynamics-aware retargeting is underexplored. This project investigates how the quality of interactively generated kinematic reference motions influences physics-based refinement and downstream motion-imitation learning. A model-agnostic interactive animation software was developed to construct, validate and export robot keyframe trajectories for the Unitree G1 humanoid and Go2 quadruped. The software has cross-platform support and is used to generate trajectories that were successfully converted and replayed in a robot-learning simulation environment. These results establish that interactive motion design can consistently produce pipeline-compatible references, providing the experimental foundation for evaluating their quality. Further work will systematically measure and compare the effects of varied

kinematically feasible motions on downstream pipelines, including optimization convergence, dynamic feasibility and tracking-controller performance in mlab. This evaluation will characterize which properties of kinematic reference most strongly affect dynamics-aware humanoid motion retargeting and provide practical guidelines for creating motions that are easier to refine, learn and deploy.

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.

Investigating oxygen vacancy formation in MnFeCoNiCuO_x high entropy spinel oxides using density functional theory

Alexis N. Lindenfelser

Mentors: Bryan Goldsmith, Kara D. Fong, Jimmy Wen, Sila Donmez, and Yu-Chuan Lin

High entropy oxides (HEOs) are an emerging class of materials with applications in catalytic conversion and energy storage. By combining five or more metal cations in equimolar amounts, HEOs have high configurational entropy that stabilizes single-phase solid oxide solutions, enabling enhanced tunability and offering distinct active sites to facilitate multistep reactions, such as the reverse water gas shift (RWGS) reaction. Experimental studies on spinel-type MnFeCoNiCuO_x HEOs found that Co- and Mn-rich samples (type 1) had enhanced RWGS activity compared to Fe-, Cu-, and Ni-rich samples (type 2), despite similar configurational entropies. We hypothesize that local mixing effects from distinct cation coordination environments influence the reversibility of oxygen vacancy formation and thus modulate RWGS activity. In this work, we use partially inverse CoFe₂O₄ spinels, generated using special quasirandom structures, as a computationally tractable structural analog for MnFeCoNiCuO_x HEOs. Mn, Ni, and Cu are substituted into local coordination shells around surface oxygen sites in (100) slab models to simulate type 1 and type 2 atomic environments. We compute oxygen vacancy formation energies using density functional theory (DFT) calculations, demonstrating that near-neighbor effects explain the activity discrepancy between samples, and showing that HEOs can be tuned by elemental compositions to improve catalytic activity.

Design of novel nanobodies targeting newly identified conserved regions of HIV gp120 and gp41 glycoproteins

Anthony K. Lindo

Mentors: Stephen L. Mayo and Andres Orta

The acquired immunodeficiency syndrome (AIDS) epidemic has been a global public health concern for over 40 years, infecting 84 million people worldwide and resulting in approximately 40 million deaths total. AIDS is caused by the retrovirus, Human Immunodeficiency Virus 1 (HIV-1), which specifically targets CD4+ T cells resulting in an immune-deficient state over a period of 2-10 years. Recently, the Mayo Lab in collaboration with the Bjorkman Lab discovered an invariant epitope within the HIV Envelope known as the Caldera Region. This presents a potential therapeutic to HIV where a specialized antibody, known as a nanobody, can target the Caldera Region and prevent host cell membrane fusion. Leveraging diffusion models for de novo protein design, this project aims to utilize Boltzgen and Boltz-2 to generate nanobody binders to target the Caldera Region and potentially

alternative invariant regions. In order to validate Boltzgen, control nanobodies were designed against the closed HIV envelope which demonstrated predictable binding in deglycosylated regions. Ongoing testing with the fully glycosylated envelope reveals candidate nanobodies targeting a functional pocket that potentially corresponds to the CD4 binding site HIV uses to integrate into CD4+ T cells.

Recovering settling velocity distributions of mud and flocs from light sensor data

Amber G. Liu

Mentors: Michael P. Lamb and Kimberly L. Miller

Fine sediment in rivers clumps into loose, low-density clusters called flocs. How quickly they settle controls where mud and organic carbon deposit, which shapes floodplains and deltas worldwide. Because flocs break apart when sampled directly, their settling must be measured using a non-invasive method. This project recovers how fast suspended sediment settles from the way it dims light as it falls in a laboratory tank. I developed a model that predicts the light signal for any set of settling speeds, and the measurement is run backward to recover them. The recovered speeds are then checked against a second, independent method. I tested the approach in stages, using single grain sizes, then mixtures, then flocs, comparing against known grain sizes and video tracking. Across all three stages, the recovered speeds match the references within a factor of two to three, and several independent solution methods agree. Flocs settle far slower than solid grains of the same size, and added organic matter turns about a quarter of the clay into flocs. The result is a validated, non-invasive way to link light measurements to fine-sediment settling and to how organic matter drives flocculation.

Sparse predict and demote: Optimal scheduling with minimal size advice

Chengyi Liu

Mentors: Adam C. Wierman and Tinashe Handina

Knowledge of job sizes is crucial for optimizing the mean response time in single-server scheduling. However, predictions of job sizes are typically inexact and may be costly to obtain. Thus, a long-standing research question is: what is the minimal information about job sizes required to optimize the mean response time? In this work, we introduce a novel policy, Sparse Predict and Demote (SP&D), and prove that it achieves order-optimal mean response time in heavy traffic while using only one-bit predictions about the size of a vanishing fraction of arriving jobs. SP&D achieves order-optimal mean response time even in situations (i) when predictions are inexact up to any fixed multiplicative error bound, and (ii) when predictions are costly, i.e., the server is used to compute predictions and is thus unavailable to serve jobs while processing prediction queries. Finally, we use numerical experiments to illustrate that the SP&D maintains near-optimal performance beyond the heavy-traffic regime.

Physics-informed 3D shape generation using Laplacian preconditioned guidance

Hudson Liu

Mentors: Anima Anandkumar and Xinyi Li

Denoising diffusion models are a leading approach to 3D shape generation tasks across multiple representations, including point clouds and signed distance fields (SDFs). A popular task for such models is to generate shapes that satisfy physics constraints. However, direct application of diffusion guidance can be unstable for shape generation, due to noisy and sharply localized gradients. We propose preconditioning the guidance update with an operator derived from the Laplace-Beltrami operator, smoothing the per-position gradients over the surface through a single implicit step of the heat equation. We evaluate our method on TIGER, a point-cloud generation model, and SDF-Diffusion, which produces truncated SDFs (T-SDFs) discretized on a regular grid. For TIGER, we define a kNN uniform Laplacian over the point cloud; for SDF-Diffusion, we define both a mollified cotangent Laplacian on the extracted zero-set mesh, and a projected Laplacian operator defined directly on the T-SDF grid itself. As guidance objectives, we introduce a bilateral-symmetry loss and an upright-stability loss.

Depressive symptoms and momentary mood relate to decision-making under uncertainty

Marcel Z. Liu

Mentors: Robb B. Rutledge and John P. O'Doherty

Everyday decisions shape important outcomes in people's lives, making it important to understand why people choose as they do. Depression affects how people evaluate potential risks and rewards, but it remains unclear whether decision-making is shaped primarily by stable symptoms or by changes in momentary mood. To examine these relationships, we replicated and extended previous analyses using a longitudinal dataset from 481 depressed patients across 4,554 eligible task sessions. In each session, participants repeatedly chose between a certain option and a gamble. Some gambles involved only possible rewards, whereas others included possible gains and losses. Participants generally preferred certainty and weighted potential losses more heavily than equivalent gains, leading them to gamble less when they stood to lose something. A broad depression component was associated with less risk-taking in trials that included possible gains and losses and with greater loss aversion, meaning that participants behaved as though potential losses were larger than equivalent gains. A second component, characterized by greater anxiety and lower anhedonia, was associated with greater risk aversion. The influence of symptoms on behavior appeared to be trait-level, as session-to-session changes in mental health and emotion did not reliably predict changes in gambling within the same person. However, higher-than-usual happiness immediately before a choice was associated with more gambling when only gains were possible. Overall, these findings suggest that stable symptom differences and immediate emotional states may relate to distinct aspects of decision-making under uncertainty.

Walk-on-Boundary supervision for linear elasticity neural operators

Siqi Liu

Mentors: Anima Anandkumar and Xi Deng

The linear elasticity partial differential equation models the deformation of structures under applied forces, making it important to engineering design. The industry standard for solving this PDE is the finite element method (FEM), but it requires a volumetric mesh, which can be expensive to construct and also introduces deterministic discretization error. Monte Carlo methods such as Walk-on-Spheres (WoS) avoid volumetric meshing, but WoS is designed for scalar equations and cannot be applied directly to the coupled, vector-valued linear elasticity PDE. Our project aims to implement a volumetric-mesh-free stochastic solver for linear elasticity and use its estimates as noisy labels for the supervised training of neural operators. Since the PDE has a boundary-integral representation, we implement a Walk-on-Boundary (WoB) algorithm using a matrix-valued elasticity kernel. A main challenge is that the WoB estimates are noisy, so we apply local importance-sampling methods, such as three-plane sampling, to reduce their variance. Next, we will generate datasets that vary geometry, material parameters, boundary conditions, and samples per point, then train a neural operator to predict solution fields from the noisy labels. This will allow us to study whether neural operators can learn accurate solution maps from inexpensive, noisy Monte Carlo supervision.

Layer selection of VLM backbone for VLA

Yuxiang Liu

Mentor: Aaron D. Ames

We conducted a comprehensive study regarding the two underexplored questions of the VLA community: (A) which layers of VLM backbone serves as the best action interface to bridge action head? (B) Does an action interface that fuses outputs of multiple layers outperforms an action interface using a single layer? To answer the questions, we performed a layer-sweep experiment based on GR00T. We iterated over all 28 layers of the COSMOS backbone and trained one VLA using each of these layers as an action interface. The performances of each layer are compared with each other and a baseline that fuses multiple layers' outputs as the action interface. We obtained two main results: (A) The performance of VLA generally grows with layer depth at the first shallowest layers,

which then enter a flat plateau with stable performances. Then, a sharp decline of performances is observed at the few deep most layers. (B) Layer fusion contributes no performance boost compared to single-layer interface. We further theoretically reason about both observations using information bottleneck.

Determining the differential progression of specific aneuploidies at the onset of gastrulation

Isabella T. Lo

Mentors: Magdalena D. Zernicka-Goetz and Ayse Damla Durmaz

Early human development demonstrates a remarkable tolerance for aneuploidy with over 80% of blastocyst embryos displaying mosaicism, and recent single-cell RNA-sequencing studies suggest that mosaic aneuploidy persists to at least week 26. Previous studies investigating mosaic aneuploidy have discovered a lineage specific depletion of aneuploidy, especially in embryonic lineages, but have primarily focused on reversine-induced aneuploidy, a system that generates chaotic, non-specific aneuploidies which are non-representative of in vivo development. However, due to ethics of culturing human embryos past day 14, the progression of aneuploidy in human development is largely unknown in gastrulation, a key developmental milestone in the black box of development. By utilizing human embryonic stem cells as a model for the epiblast lineage, we study the progression of aneuploidy through early gastrulation. Here, we compared trisomy 12 and trisomy 21 cell lines, reversine-treated cells, and mixed euploid:aneuploid chimeras in the lumenoid model. By staining for the EMT marker Brachyury to assess progression through the primitive streak, we found trisomy 12 lumenoids showed reduced brachyury expression when compared to trisomy 21, reversine-treated, and euploid lumenoids. Furthermore, chimeric conditions tended to show a decreased proportion of brachyury-positive cells within each lumenoid.

Characterizing estrogen-related receptor function during nephron tubule development using zebrafish

Kara Z. Lo

Mentors: Marianne Bronner and Tianli Qin

The nephron, the kidney's fundamental structural and functional unit, performs essential physiological functions, including metabolic waste excretion, hormone production, and maintenance of fluid and electrolyte homeostasis. Each nephron consists of a blood-filtering unit followed by an epithelial tubule patterned into functionally specialized segments. Proper nephron development requires coordinated segment patterning and tubule morphogenesis to establish a functional renal tubule, and defects in these processes contribute to kidney disease. We use the zebrafish embryonic kidney, which exhibits a high degree of genetic and cellular conservation with the mammalian nephron, as a high-throughput model to investigate the roles of the estrogen-related receptors *Esrra* and *Esrrb* in nephrogenesis. Analysis of CRISPR-Cas9-generated single and double knockout embryos using hybridization chain reaction (HCR) reveals expansion of the proximal convoluted tubule (PCT) and concomitant shortening of the distal tubule, suggesting impaired tubule convolution. Because proper tubule convolution depends on cilia-driven fluid flow that directs collective migration of nephron epithelial cells toward the proximal end, as well as stretch-induced proliferation in the distal tubule that sustains this process, ongoing studies are testing whether these phenotypes arise from defects in cilia function or reduced cell proliferation.

Vertices of the convex polytope consisting of all monotonic stochastic choice functions

Vivian Loh

Mentor: Kota Saito

In economics, empirical data on individual or aggregate choice behavior frequently takes the form of stochastic choice data. Market share data, for instance, records the fraction of a population selecting each alternative from a given menu. Define a *stochastic choice function*, from the power set 2^X of the set of alternatives X to $\mathbb{R}$, to be one that, for each menu $S \subseteq X$, assigns a non-negative probability to each element in S with total probability 1 for that subset. The convex polytope consisting of all linear combinations of *linear orders*, precisely defines the random utility models (RUM) -- the functions which satisfy the property that their Block-Marschak polynomials are nonnegative. If we consider

monotonicity, the property that adding an alternative into the menu does not increase the probabilities of the other alternatives, then all monotonic choice functions form another convex polytope. We compute this convex polytope for small values of $|X|$, and suggest models to characterize some or all of its vertices. Eventually, we completely characterize the subset of vertices that have Stochastic Sequential Criteria (SSC) representations, and compute the number of such vertices for each number of total alternatives. Efficiently characterizing the monotonic choice functions will simplify future work in behavioral economics by limiting the space of studied choice functions to only the relevant ones.

Finding periods in variable stars using deep learning and dm-dt images

Tejas A. Lohia

Mentors: Ashish Mahabal and Matthew J. Graham

Estimation of periods from stellar light curves is an essential task in time-domain astronomy, as it enables the characterization of variable stars and provides valuable insights into stellar structure and evolution. Traditional period-finding algorithms face several challenges, including irregular sampling, measurement uncertainties, and data sparsity. In this project, we focus on building deep learning models to regress the periods of variable stars. We transform the light curves into dm-dt images that capture pairwise differences in magnitude and time, enabling the use of image-based deep learning models. The models are being trained and evaluated on synthetic data and observations from the TESS and ZTF surveys, allowing us to assess their period estimation accuracy and ability to generalize across surveys with different observing cadences. Current work focuses on improving cross-survey generalization by generating cadence-faithful synthetic light curves that better resemble real survey observations. This work is intended to enable accurate, automated period recovery for variable stars across current and upcoming time-domain surveys, thereby supporting downstream classification and stellar characterization efforts.

Continuous best-response learning for direct exploitability descent

Eric Y. Lu

Mentor: Tony Yue Yu

Directly minimizing exploitability provides a principled approach to approximating Nash equilibria. However, this requires repeatedly computing or learning best-responses, which are not necessarily continuous with respect to the opponent strategy profile. Therefore, we propose Deep Q-Network Exploitability Descent (DQN-ED), an algorithm that learns continuous optimal state-action values instead. We evaluate this method in three-player, three-card, imperfect-information Goofspiel. Under matched environment transitions and comparable learning configurations, across 5 random seeds, DQN-ED reaches a final mean NashConv of 0.0046, compared to the 0.0857 baseline set by Neural Fictitious Self Play.

Dose-controlled In-Ga eutectic delivery for oxide destabilization in additively manufactured ceramic structures

Mark N. Luo

Mentors: Katherine T. Faber and Zachary Chase

Laser powder bed fusion (L-PBF) enables fabrication of architected Al-Al₂O₃-YSZ ceramic composites. Still, the native oxide skin on aluminum particles inhibits wetting, causing molten aluminum to pool rather than infiltrate the pore network during reaction bonding. A liquid Indium-Gallium (In-Ga) eutectic destabilizes this oxide skin. It improves wetting within a critical thermal window, but single-droplet delivery is unreliable on non-flat geometries because the eutectic remains a mobile droplet at room temperature. In this project, a geometric-independent dose-controlled In-Ga strategy was developed. After an initial approach using a porous ceramic sleeve with alumina powder proved difficult, work shifted to surface-coating methods: aluminum-powder pre-layering, In-Ga-coated weighing-paper wraps, and direct brush application tested against single-droplet dosing on staircase-cut coupons. Direct brushing, leveraging the eutectic's strong self-cohesion, produced the most complete surface coverage; coating quality depended on surface finish and moisture, while heating showed no clear effect. Electron probe microanalysis and wavelength-dispersive spectroscopy of staircase coupons revealed reduced indium concentration in regions blocked from direct application,

while gallium remained unaffected, suggesting the two elements may not migrate together. Remaining work will confirm this transport behavior and compare coating methods to establish a scalable, geometry-independent In-Ga infiltration framework.

Biosynthesis and cellular effects of phenylacetyl-carnitine, a gut microbiota-derived metabolite linked to Alzheimer's disease

Xingchuan Ma

Mentors: Sarkis K. Mazmanian and Daniel Dumitrescu

Variations in gut microbiome composition have been observed in humans with Alzheimer's disease (AD), yet whether microbially-produced small molecules directly contribute to AD pathophysiology is poorly understood. Recently, untargeted metabolomics profiling of AD mouse models discovered a novel microbial metabolite, phenylacetyl-carnitine (PAC), which was significantly enriched in the serum of AD mice and correlated with cognitive impairment in human AD cohorts. However, how PAC is biosynthesized and what effects it has on host cell physiology is unknown. Whereas PAC production was not detected by complex gut microbiome communities under anaerobic conditions, I found that HepG2 liver cells incubated with phenylacetate (PAA) and carnitine actively synthesized and secreted PAC into the culture medium. Notably, PAC biosynthesis was significantly decreased upon co-treatment with inhibitors of carnitine palmitoyltransferase 1A (CPT1A) and carnitine O-acetyltransferase (CrAT). I also found that PAC-treated HepG2 cells exhibited decreased proliferation and increased levels of reactive oxygen species (ROS). Given the importance of acylcarnitines in lipid metabolism and energy homeostasis, I am using Seahorse assays to measure mitochondrial respiration in cells treated with or without PAC. Altogether, this work explores PAC as a candidate molecular mechanism linking gut microbial metabolism and mitochondrial dysfunction in AD.

Building native-resolution maps and interactive visualization for Enceladus's south polar terrain

Yingying Ma

Mentor: Mark Simons

Understanding Enceladus's south polar terrain requires maps that keep the resolution and coordinate system of each source clear. This project develops a mapping and visualization workflow based on the version-2 Enceladus shape model. The workflow combines regional grids sampled at 25, 50, and 100 m only within their valid coverage. Finer grids take priority where they exist, and companion resolution maps record the source used at each location. Current outputs include local, medium, and large north-up/east-right relief maps and a QGIS workspace. The QGIS workspace overlays the available mapped feature layers on the topography for comparison. The project also includes geodesic profiles and an interactive tool that links a rotatable globe, a two-dimensional map, and raw topographic cells from each profile swath. These products provide a consistent way to inspect the target and prepare figures. Current work on the Cassini image archive covers file organization and metadata checks; image-to-topography registration has not yet begun.

[4+2] Cycloaddition on β -ketoester with tethered alkene and siloxane tether

Sarah E. Madden

Mentors: Brian M. Stoltz and Sara Siddiqui

Metal-catalyzed allylic alkylation is a well-explored process that allows for the synthesis of organic molecules. In particular, variations of decarboxylative asymmetric allylic alkylation have been studied extensively on β -ketoesters over the last several years. Protonation, conjugate addition, and the aldol reaction have been performed on variations of a chiral palladium enolate already, as well as a [4+2] cycloaddition. In order to access a more diverse variety of products and scaffolds with a palladium enolate undergoing a [4+2] cycloaddition, we are modifying the β -ketoester substrate commonly used by the Stoltz lab. We are appending an alkyne tethered to the enone through a siloxane linker, which will serve as a synthetic handle for further product manipulation. The siloxane will likely break the five-membered carbocycle that has formed on past attempts of a [4+2] cycloaddition on similar substrates within the lab. Then, we will test and refine the conditions of the [4+2] cycloaddition on this substrate.

Advancing hybrid mRNA vaccine strategies through the development of next-generation ESCRT-recruiting domains

Indeever Madireddy

Mentors: Magnus Hoffmann, Pamela J. Bjorkman, and Kim-Marie Dam

Abstract withheld from publication at mentor's request.

Development of a standalone scintillation detector kit for muon and gamma detection

Samikshya Mahapatra

Mentor: Ryan B. Patterson

Scintillation detectors are one of the most fundamental particle detectors and are essential for scientific discovery, but the few existing educational kits conceal their electronic components and the physics behind detection, preventing students from fully understanding the detection process. Here, we develop an easily reproducible, standalone scintillation detector kit to help students actively engage with particle detection principles and circuit design. The system uses easily sourced plastic or thallium-doped cesium iodide scintillators coupled with silicon photomultipliers, powered by USB power supplies. Throughout the development process, various amplifier, comparator, and counter configurations were designed and tested to optimize the signal readout. To verify particle interactions, particularly from cosmic muons, a stacked scintillator arrangement was implemented to accurately measure coincidence signals. Additionally, noise reduction techniques, such as RC filters and electromagnetic shielding, were integrated to improve the signal-to-noise ratio. Results demonstrate successful amplification of silicon photomultiplier signals and reliable detection of muon coincidence events. Two versions of the system have been developed based on the energy range to be studied. These results confirm the viability of a transparent and accessible detector design, providing a practical educational kit that allows students to construct their own independent particle detector to measure naturally occurring radiation.

Vision-guided precision targeting for spatially controlled biolistic corneal delivery

Maryam Y. Malkosh

Mentors: Julia A. Kornfield and Dennis L. Ko

Millions of people worldwide experience visual impairment caused by diabetes mellitus and delayed corneal wound healing, highlighting the need for accessible, low-risk, and cost-efficient treatment strategies. The Kornfield lab has developed a biolistic delivery platform for targeted administration of therapeutic particles to the cornea. In this project, we integrated the delivery device with a robotic positioning stage, image processing, feature detection, and tracking to enable spatially controlled particle delivery with approximately 2 mm x 2 mm localization. Specifically, we tested our vision-guided targeting system with 2D images, 2D images embedded within a gel, and *ex vivo* porcine eyes. Particle delivery positions were validated with optical coherence tomography and overhead camera images and processed with Fiji and Python. With this system, the user is able to target specific ocular regions, allowing for potential applications targeting the limbus with regenerative therapeutics to promote corneal epithelial repair. Future work will focus on optimizing the system for clinical translation, including integrating targeting with corneal topography and real-time motion adaptations for *in vivo* deliveries.

Expanding the antigen repertoire and characterizing localization dependence of programmable nanobody circuits for selection

Neev Mangal

Mentors: Katie Galloway, Rustem F. Ismagilov, and Mary Ehmann

Cell purification methods are a major bottleneck in advancing cell-based therapeutics. Conventional separation methods rely on flow-based or physical separation processes which reduce cell viability and scale poorly with cell numbers. DASIT (Destabilized-nanobody Antigen Selection and Identification Tool) operates at the protein level, fusing a destabilized nanobody to an antibiotic resistance protein to selectively provide antibiotic resistance to antigen-positive cells. While DASIT functions robustly with EGFP and ALFA nanobodies, not all nanobodies tolerate destabilization or domain fusion. Here, we

evaluated FLAG and BC2 nanobodies, showing that wildtype nanobody constructs survived puromycin selection without enriching for antigen-positive cells. Additionally, both wildtype and destabilized BC2 constructs lost mRuby2 signal after lentiviral delivery despite apparent integration, implicating a transcript-level defect. We also evaluated DASIT's selection ability for nuclear-localized antigens in HEK293T cells. Enrichment was moderate in cases when both the antigen and the DASIT construct were permanently co-localized, weakest when only the antigen was nuclear-localized, and strongest (85% EGFP+) when only DASIT was nuclear-localized. This result suggests that a post-translational binding period in the cytoplasm is critical to DASIT's success. In hiPSCs, we discovered that modRNA-delivered DASITEGFP-PURO enriched LAMP1-EGFP hiPSCs cells from ratios as low as 1:14 (LAMP to wt iPSCs) but failed to enrich SOX2-EGFP hiPSCs. Overall, these results show DASIT can be extended to EGFP antigens with different cellular localizations (lysosome, nucleus), but further work is needed to expand DASIT to other antigens and to generate rules for reproducible nanobody destabilization.

Assessing impact of math camp active-learning activities on students' confidence and perceptions towards STEM

Marclus M. Marite

Mentors: Franca Hoffmann, Jasmine R. Bryant, and Joe Champion

Mathematics camps offer an alternative learning environment emphasizing collaboration, active participation, and inquiry-based learning. This project evaluates how active-learning mathematics camps affect students' mathematical confidence, attitudes toward mathematics, and interest in STEM, using a mixed-methods design combining archival analysis of student reflections, interviews with former camp volunteers, and questionnaires administered to current participants and alumni.

Since the first interim report, the 2026 camp ran from 10–14 August in Nyarongi, Homa Bay County, serving approximately 50 high-school and 60 pre-school/primary-school students, supported by roughly 15 volunteers. Entry and exit questionnaires were administered to student participants (33 and 42 responses respectively), and semi-structured interviews with former volunteers have been conducted. Baseline attitudes toward mathematics were already high before camp began, consistent with more than half of entry respondents having attended a previous Nyarongi camp; a preliminary comparison shows broadly stable self-reported attitudes with no large shifts, alongside evidence of ceiling effects to be examined further once qualitative data are integrated. Remaining work includes transcribing and coding the historical journal archive, coding volunteer interview transcripts, and distributing questionnaires to alumni and first-time volunteers.

Design and integration of deployment mechanisms for ultralight rollable booms and an elastically-foldable structure

Mahak Mathur

Mentors: Sergio Pellegrino and Flavia Palmeri

The Space Solar Power Project (SSPP) aims to provide continuous solar power to Earth by wirelessly transmitting energy absorbed from the Sun by large in-space, autonomously deploying solar arrays. The Deployable on-Orbit ultraLight Composite Experiment (DOLCE) was the first SSPP demonstration mission. DOLCE consists of an elastically-foldable solar array structure, four supporting ultralight rollable booms, and their respective deployers. The square structure, composed of four identical triangular sectors, is packaged by folding each sector into a flat strip and coiling it around the main-stage deployer. During DOLCE deployment, the booms and structure were deployed in a decoupled sequence, increasing system complexity and the risk of misalignment and twisting. To address these issues, a new deployment scheme was developed. In this approach, the main-stage deployer uncoils the structure while the supporting booms, connected to the structure through diagonal tapes, are synchronously unrolled by the boom deployers. Upon uncoiling, the booms extend and tension the tapes, guiding the unfolding of the structure. Implementing this deployment scheme requires the integration and synchronization of the main-stage and boom deployment mechanisms. The DOLCE

deployment mechanisms were thus modified and integrated into a single cohesive system. Several prototype iterations were developed to demonstrate the synchronized coiling/uncoiling of the structure with the boom retraction/deployment, as well as the controlled unfolding of the structure, while minimizing volume and mechanical complexity.

Arrival-time statistics of a searcher in a cue field: First passage, resetting, and capture

Philip-David Medows

Mentors: Rob B. Phillips, Minakshi Ashok, and Sara Mahdavi

Many living things must find a target they cannot see, from a bacterium swimming toward food to an immune cell following a chemical trail, guided only by a local sense of whether conditions are improving. We ask a basic quantitative question about such searches: how long do they take? Modeling the searcher as biased diffusion through a spatial cue that shapes its motion, we derive not merely the average time to reach an absorbing target but all the moments of the full distribution of arrival times for a general cue field, with the constant-drift case as an exactly solvable benchmark. We then extend this first-passage description in three directions. Under stochastic resetting, a searcher that restarts at random intervals always arrives eventually, and there is a genuine best restart rate, which helps precisely when the search is more variable than it is long. Under partial absorption, in which the target binds only sometimes on contact, we obtain the mean capture time in closed form. And we show that this capture time separates exactly into the time spent traveling to the target, which depends on where the search began, plus a single reaction time that does not. Together these results give a nearly closed-form account of when a guided searcher reaches, and is captured by, a target it never sees.

What heads are you really finding? Single-token vs. free-form localization

Ayushi Mehrotra

Mentors: Dylan Hadfield-Menell, Adam C. Wierman, and Aruna Sankaranarayanan

Task settings that elicit a single-token response from a Language Model are widely used in interpretability research to causally localize the circuits underlying the task. Yet, this formulation is poorly suited to complex behaviors such as factual recall, reasoning, summarization, and style transfer. Prior work therefore often studies these behaviors through multiple-choice question answering (MCQA) formulations, or leaves them unexplored altogether. In these constrained settings, a model is said to exhibit the target behavior when it predicts the aligned answer choice. We propose, and defend, an alternative: localizing behavioral circuits directly from free-form generative responses. Across five model families and four tasks with matched MCQA and generative variants, we find that circuits localized from free-form responses are more precise than their MCQA-derived counterparts. We employ mechanistic analyses to explain this difference: MCQA introduces intermediate computations that are absent during direct generation, and contains high-rank representations. Taken together, our results suggest that localizing circuits from free-form responses provides a more direct and effective basis for model control

Enhancing predictive generative policies with learned rewards

Divan I. Mejia Gonzalez

Mentors: Yilun Du, Adam C. Wierman, and Haldun Balim

Many manipulation tasks give a reward only on completion, and writing a denser reward by hand is difficult and specific to each task. Reinforcement learning algorithms for learning critics have been shown to be brittle in this regime. We design a planning algorithm for generative policies that instead uses a reward learned directly from expert demonstrations. This reward, which we call the progression function, defines a preference over states through the Bradley Terry model, and is trained to favor states closer to task completion in the demonstrations. The generative prior and the reward are combined within a model predictive control framework. At each decision the policy proposes a set of candidate action sequences, a learned world model predicts the states each would reach, and the

progression function scores those predicted states, yielding an objective for optimization. We evaluate our method on manipulation tasks from the robomimic benchmark, with both low dimensional state and camera observations, and demonstrate improvements over the base policy and over our algorithm when using other learned rewards.

Has Europa always hosted an ocean? Coupled thermal-orbital modeling of a hydrosphere grown by mantle dehydration

Aliaksandr Melnichenka

Mentors: James W. Fuller and Amir Bagheri

Europa hosts a global liquid water ocean beneath its ice shell, and the age of that ocean bears on how long the moon has been habitable. Almost every previous model has assumed the ocean is primordial, because the water inventory is fixed as a constant before the integration begins. The alternative model is that Europa accreted essentially dry and released its water later, as hydrated silicates in the rocky mantle warmed and broke down. This project rebuilt a coupled thermal and orbital model of the Laplace resonance so that the hydrosphere is grown by the thermal history instead of assumed. Water leaves the rock once the mantle crosses the dehydration temperature, mass moves from rock to hydrosphere at fixed total mass, the ice shell thickness follows from the conductive heat flux, and the tidal response and the moment of inertia track the evolving structure. About 8000 histories spanning 4.5 Gyr were integrated from starts with no ocean and no ice, sampling the initial mantle temperature, the extractable water fraction, the radiogenic abundance, the core density and the dehydration timescale. Each was scored against five present-day observations, namely a liquid ocean, the conductive ice shell thickness, the total thickness of ice plus ocean, the surface heat flux, and the measured moment of inertia. Europa can build its ocean out of its own rock. From a dry start, radiogenic and tidal heating drive dehydration until an ocean appears, and in some histories that satisfy every constraint it appears more than a billion years late. Chondritic radiogenic abundances alone never warm the mantle enough, so a late ocean requires an enhanced inventory of uranium, thorium and potassium. That same requirement bounds the delay, since a mantle rich enough in isotopes to melt late ends up too hot and too differentiated to match the measured moment of inertia. No accepted history forms its ocean more than about 2 Gyr after the moon, which places a floor of roughly 2.5 Gyr on the age of the ocean. These results are preliminary and rest on a simplified linear dehydration law.

Theoretical optimization of pick-up coil configurations for vibrating sample magnetometry

Christopher S. Meza

Mentors: Axel Scherer and Mladen Barbic

Accurately measuring the magnetization of materials relies on maximizing the induced voltage in the pick-up coils of a vibrating sample magnetometer. Historically, established coil geometries have provided a baseline for these measurements, but advanced materials analysis requires greater sensitivity and lower noise. This project investigates a new, optimal pick-up coil configuration and ideal vibration direction to maximize signal detection. Using numerical methods and computational simulations, we evaluated the spatial distribution of magnetic field derivatives to dictate precise coil positioning and shape. Our analysis identified novel coil geometries that yield at least two and a half times the signal pickup of standard configurations. These specially shaped and oriented coils capture significantly more of the induced signal. Implementing this optimal configuration will enable more accurate readings of magnetic moments, substantially reducing noise and improving the precision of future magnetic measurements.

Masked autoencoder pretraining for stable PDE rollouts

Annabelle Miin

Mentors: Anima Anandkumar and Robert Joseph George

Neural operators can predict the next state of a physical system accurately, but small errors compound when the model is rolled forward over many steps. We ask whether masked autoencoder (MAE) pretraining, teaching the model to reconstruct corrupted inputs before any supervised training, can make these rollouts more stable. Prior work has applied MAE to PDE data using vision

transformers and spatial patch masking, a corruption borrowed from vision. However, PDE information lives in frequency content, so we corrupt in the frequency domain instead, removing low, high, band, or random coefficient content in Fourier or wavelet space. We pretrain a MAE with a Fourier Neural Operator (FNO) backbone on two datasets from The Well benchmark with each strategy, finetune it for forecasting, and compare against a supervised baseline given the same total compute. Our results show that MAE pretraining in the frequency domain improves mean rollout by 40% on turbulent radiative layer and 12% on active matter compared to baseline.

Design of a quantum bus for simultaneous all-to-all connectivity between superconducting qubits

Todor Miljevic

Mentors: Mohammad Mirhosseini, Frank Yang, and Parth Shah

Currently, superconducting quantum processors are inherently bottlenecked by quantum decoherence and connectivity limitations to nearest-neighbor only. Although quantum error correction combats decoherence, a far too large of a hardware overhead is required to actually implement it. The lab attempts to overcome both of the aforementioned bottlenecks by designing modular chips in such a way that simultaneous all-to-all connectivity is possible within each module. The lab demonstrates computationally an all-to-all modular qubit connection system using a resonator array scheme, which enables communication through evenly-spaced normal-mode frequencies centred around 5GHz. Along with demonstrating the array's functionality using an electromagnetic solver, we predict its robustness to noise to be sufficient for tolerating up to 7MHz of frequency disorder per resonator.

Mathematically modelling neuronal circuit recovery through activity dependent myelination

Elizabeth J. Milne

Mentor: Michelle Effros

In "Unsupervised restoration of a complex learned behavior after large-scale neuronal perturbation", a paper by the Carlos Lois group, it was found that after genetic perturbation of the majority of excitatory neurons in the HVC, a brain region involved in song production, of zebra finches, the bird's song severely degraded, then fully recovered within two weeks. The mechanism behind this recovery is not completely understood; some network-level recovery mechanisms are proposed in the Lois group paper. In this report, we investigate the feasibility of potential recovery mechanisms operating within a very small neuronal circuit (involving 3 neurons), in particular the potential role of activity-dependent myelination. To do this, we utilize a mathematical framework developed by Effros and Li, which uses Markov chains to analyze the long-term behavior of neurons. By analyzing the long term behavior of the post-synaptic neuron with different statistical relationships between the two presynaptic inputs, we seek to find the conditions under which an unmanipulated neuron can adapt, via activity-dependent myelination, to compensate for a silenced neuron.

Efficient entanglement-free benchmarking of logical magic state: The suggestion of protocol and theoretical boundaries

Mao Mishima

Mentors: John P. Preskill and Senrui Chen

Magic states enable the non-Clifford operations required for universal fault-tolerant quantum computation. Since preparing high-quality magic states is costly, it is desirable to benchmark them using few copies and without additional magic resources.

We study relative infidelity estimation using stabilizer measurements that act on each copy separately, without entangling different copies. For a fixed number of qubits and a general known pure target, we derive a structural criterion that determines when such estimation is sample-efficient. Specifically, the required number of copies scales as either the inverse infidelity or its inverse square. This distinction is particularly important for fault-tolerant applications, where very low infidelities are required and the difference in sampling cost becomes substantial.

For computational-basis phase states with equal-magnitude amplitudes and phases given by roots of unity of power-of-two, this criterion reduces to a simple condition on the phase function, giving a complete classification within this family. We also construct explicit protocols for the efficient phase-state regime using only CNOT gates, Pauli measurements, and classical feed-forward. In particular, CCZ states and more general hypergraph states admit inverse-infidelity scaling, whereas tensor-product T states require inverse-square scaling.

Colored Jones tails, stability series, and the Gukov-Manolescu series

Arpit Mittal

Mentors: Sergei G. Gukov, Josef Svoboda, and Lara San Martin

The colored Jones polynomials of a knot are a family of Laurent polynomials associated with irreducible representations of $U_q(\mathfrak{sl}_2)$. For positive and alternating knots, the limit of the colored Jones polynomials divided by the lowest term converges to a q -series, known as the tail of the knot. Garoufalidis and Lê extended the definition of the tail to a stability series, which they proved exists for alternating knots. Based on some examples, Park proposed a possible definition for the Gukov-Manolescu series for strongly quasi-positive (SQP) knots, which was later conjectured to be equivalent to the stability series. We study the existence of the stability series for positive knots. We prove that the tail of a Murasugi sum along the canonical Seifert surfaces of positive knots is the product of their individual tails. Additionally, we disprove a conjecture by Orland et al. by identifying an SQP knot with no tail. We also study the asymptotics of some tails and find some new relations, such as between the tail of 8_15 and a q -series studied by Garoufalidis and Zagier, which is related to the Kashaev invariant of 4_1 .

Simulating and tracking domain evolution in the quantum Ising ferromagnet $\text{LiHo}_{0.65}\text{Y}_{0.35}\text{F}_4$

Umar H. Mohammad

Mentors: Thomas F. Rosenbaum and Daniel Silevitch

Magnetization reversal in a ferromagnet occurs through discrete jumps in an ensemble of domains, which, when detected with an inductive pickup coil, produce audible signals known as Barkhausen noise. At low temperatures, ferromagnets in the $\text{LiHo}_x\text{Y}_{1-x}\text{F}_4$ family can exhibit quantum-tunneling-driven domain dynamics, producing "quantum Barkhausen noise". However, experiments can only measure bulk properties, limiting the resolution at which they can describe the behavior of individual domains. To study these microscopic effects, I simulate the hysteresis of $\text{LiHo}_x\text{Y}_{1-x}\text{F}_4$ using a Metropolis Monte Carlo method with extended acceptance criteria, adding a transverse-field-dependent tunneling term in parallel with the standard thermal term. I developed algorithms to identify domains, track their evolution by maximizing intersection-over-union between frames, and identify domain-level Barkhausen events. I found that the domain count is bimodal during magnetization reversal and that domain-level Barkhausen events follow a power law at the domain level, mirroring the experimentally observed behavior.

Categorical Lyapunov theory for nondeterministic dynamical systems

Piash Mondol

Mentors: Aaron D. Ames and Joe Moeller

We develop a categorical formulation of Lyapunov stability for nondeterministic discrete-time dynamical systems. Nondeterministic evolution is modeled in the Kleisli category of the powerset monad with iteration induced by Kleisli composition, while worst-case measurements are interpreted through ordered comparison objects and join aggregation. We establish Lyapunov stability criteria and study a converse construction in which stability canonically determines suitable comparison bounds.

Holography in the parafermion SYK

Madeline H. Mooney

Mentors: Maria Spiropulu and Joseph D. Lykken

We study a generalization of the Sachdev--Ye--Kitaev (SYK) model. Replacing Majorana fermions with parafermions, fractional statistics and a global D_n symmetry are introduced whilst preserving the SYK's all-to-all random interactions. As the Majorana SYK possesses a Jackiw--Teitelboim (JT) gravity dual, we investigate whether analogous holographic behavior is retained in the parafermion SYK. Presenting two models, we show one to lack a spin-glass phase---an important ansatz in holography. For the other model, we take the double-scaled limit---here reformulating traditional chord diagrams into a gas ensemble---and locally derive a Schwarzian action. The chord gas also admits (regardless of any specific model) a $\text{Diff}_+(S^1)/\text{PSL}(2, \mathbb{R})$ structure consistent with the boundary configuration space of JT gravity.

Understanding the role of astrocytes in regulating zebrafish sleep

Rachel G. Moreno

Mentors: David Prober and Olivia Eliopoulos

Sleep homeostasis and sleep disorders play a vital role in health, though the underlying biological mechanisms of sleep are not yet fully understood, particularly the role of astrocytes in regulating sleep remains unclear. To further understand the role of astrocytes in sleep, we tested whether activating zebrafish astrocytes results in increased sleep. We utilized transgenic *gfap:optoA-1AR* zebrafish that when activated via blue light optogenetic stimulation leads to rapid elevations of intracellular calcium in astrocytes. Additionally, we utilized the transgenic *gfap:iβARK* line of zebrafish to test whether loss-of-function perturbations by blocking intracellular release of calcium in astrocytes will result in decreased sleep. To test these aims, 4-day old larval zebrafish were placed in each well of a 96-well plate and behavior was recorded via infrared camera in video tracker sleep assays. The fish were then genotyped by PCR and gel electrophoresis and then analyzed via MATLAB. Preliminary results demonstrate that optogenetic activation of astrocytes results in increased sleep in comparison to wildtype (WT) fish. Additionally, we found the *gfap:iβARK* zebrafish had a decrease in sleep in comparison to the WT fish. These findings suggest that decrease of astrocyte calcium levels correlates with a decrease in sleep.

Novel method for the in-space assembly of membrane structures with a stationary robot

Blake T. Morris

Mentors: Sergio Pellegrino and Tara Venkatadri

A new in-space assembly schematic has been developed utilizing a robot which ejects bays sequentially from a stationary truss-builder to form a ring-like perimeter truss, which is broken into segments called bays that can be modeled from a top view as sides of a polygon, supporting an interior membrane structure. A demonstration of this method with the DARPA NOM4D project uses a single large, prefabricated membrane to cover the whole ring, but studies have shown that this is not feasible as structures are scaled to greater number of sides and thus, novel methods of assembly of membrane structures must be developed. A new assembly schematic was studied which utilizes several rectangular membranes layered on top of one another which connect diametrically opposing bays of the structure and are attached as bays are ejected. The optimal layering of these membranes was determined and their shape optimized to maximize coverage of the interior of the polygonal ring whilst minimizing overlap. This new class of structures generated with this assembly method was assessed for rigidity and mechanical behavior with matrix analysis. A device was built which demonstrates this assembly method. This work will directly contribute to future studies of in-space assembly, specifically in the construction of large solar arrays.

Synthetic noncommutative motives for algebraic K-theory

Samuel J. Moss

Mentors: Denis-Charles Cisinski and Matthias Flach

Synthetic category theory is a modern axiomatic framework for higher category theory independent of set theory. In this paper, we use the language of synthetic category theory to reprove and expand upon several known results from the theory of noncommutative motives for algebraic K-theory. An additive invariant is a functor from the category of idempotent-complete stable categories which preserves filtered colimits and sends split-exact sequences to products. We give a synthetic construction of a stable presentable category of noncommutative additive motives and a universal additive invariant taking values in this category. Furthermore, we show this category of additive motives is the sifted completion of a certain small category, and we show that connective algebraic K-theory is corepresentable in this category. A localizing invariant is a functor from idempotent-complete stable categories which preserves filtered colimits and sends exact sequences to pullback squares. We show that localizing invariants valued in anima form a stable category. We give a synthetic construction of a stable presentable category of noncommutative localizing motives in which nonconnective K-theory is corepresentable, and we define a universal localizing invariant valued in this category.

Study of signal response of SMSPDs exposed to high energy ionizing radiation

Cecilia I. Moxley

Mentors: Maria Spiropulu and Soham Bhattacharya

Superconductive nanowire single photon detectors (SNSPDs) are used to detect photons; they have very low noise and fast signals, which can be useful for precision timing detectors for high energy particles. However the exact mechanisms triggering the detection on the atomic scale have not been proven yet for high energy charged particles. The SNSPD is simulated in Geant4 to obtain data on the interactions. The simulation models a 3nm layer of WSi(the wire), a 0.2mm layer of SiO₂ substrate and a 0.3mm layer of aSi substrate. A 120 GeV proton is shot at the WSi. Our preliminary results suggest the proton creates cerenkov photons in the substrates. The different variations of optical properties show that 20%-60% of the simulations will have photons reflected to the WSi and absorbed there, depositing 1-6eV. The proton deposits 2.27eV in the wire 99% of the time, suggesting that it alone is not the source of the signal. This project studies the sensitivity of the photon deposits to different properties of the materials which can differ at cryogenic temperatures, in comparison to reported values in literature at higher temperatures.

Independent orbit determination for Jupiter Trojans from Gaia astrometry as a foundation for a distant-perturber search

Isabella H. Mueller

Mentor: Michael E. Brown

Batygin & Brown (2016) proposed an undiscovered giant planet (known as Planet Nine) on a wide, eccentric orbit to explain the clustered orbits of distant trans-Neptunian objects. Such a body would exert a small differential acceleration across the Solar System, detectable in principle as a coherent signal in the orbits of many small bodies. Jupiter Trojans are an attractive probe because they span a baseline of roughly 10 au. Rice & Laughlin (2019) argued that precise positions for more than about 225 Trojans would separate a Planet Nine signature from that of the Kuiper belt, while Gomes et al. (2023) forecast the constraint from roughly 8000 Trojans observed by LSST through 2035. Establishing what existing astrometry can already constrain sets a baseline for these forecasts and exercises the methods they will require.

The Gaia Focused Product Release (FPR) already delivers Solar System astrometry on a single homogeneous reference frame, with CCD level positions accurate to better than one milliarcsecond along scan at the bright end. This astrometry spans an arc of 66 months, covering 0.46 of the 11.86 year Trojan orbital period. Gaia DR4 will refine the calibration but will not extend this arc. Gaia DR5, expected around 2030 and drawing on the full 10.5 year mission, will extend the arc, reaching 0.89 of a Trojan period.

We have built an orbit determination pipeline on OrbFit that uses FPR astrometry under the INPOP19a ephemeris, weights each observation by the delivered correlated covariance, and matches the FPR observation model, timescale, and coordinate conventions. For 672 Trojans, our semi-major axes match the published FPR values to a median fractional difference of 2.43×10^{-8} (maximum 6.04×10^{-8}), below the 10^{-7} threshold David et al. (2023) use to flag anomalous solutions. Formal uncertainties agree to within factors of 0.5 and 3.0. This validated baseline supports a grid search refitting orbits under a distant perturber's tidal acceleration, over a grid of perturbation strength and direction. Holman & Payne (2016) applied this approach to Pluto and other trans-Neptunian objects, finding that a tidal approximation reproduces a full moving-planet model across the sky. Sensitivity will be expressed as a limit on perturber mass at a given heliocentric distance, for direct comparison with constraints from planetary tracking.

Spectator-error-free two-qubit gates for fluxonium qubits using a dual-mode coupler

Shuhul Mujoo

Mentors: William D. Oliver, David Hsieh, and Miguel Moreira

Superconducting quantum computers promise to solve complex problems, but they require highly accurate operations on pairs of quantum bits (qubits). The fluxonium is a promising qubit due to its long coherence times, yet building fast, high-quality two-qubit operations for it remains difficult. A leading approach links qubits through an intermediate circuit element called a coupler. However, these operations frequently suffer from spectator errors, meaning they are disturbed by the states of nearby idle qubits. This interference ultimately limits how large a processor can grow. This project develops a two-qubit operation that eliminates this error using a coupler with two adjustable modes. The second mode provides the precise control needed to suppress disturbances from neighboring qubits while preserving the main operation. To implement this, we simulate the coupled circuit to identify specific settings where unwanted coupling cancels. Experimentally, we built specialized readout and control tools: a measurement that reads three qubit states simultaneously to flag leakage errors, and a calibration that corrects magnetic-flux crosstalk to steer the coupler reliably. Ongoing work focuses on calibrating the operation and measuring how effectively it suppresses spectator errors.

Observer phases in dS_3 quantum gravity

Damian R. Musk

Mentors: Joseph D. Lykken and Maria Spiropulu

We study complex phases in the semiclassical partition function of three-dimensional de Sitter gravity. We first derive the one-loop graviton contribution on the three-sphere and on lens spaces, retaining all effects from gauge fixing, conformal and zero modes, conformal Killing vectors, contour choices, and normalization. We then reassess the proposed equivalence between metric and Chern-Simons formulations, with particular attention to overlooked phase and measure factors. Finally, we translate recent observer-based interpretations of gravitational phases into Chern-Simons language, testing whether observers are represented by Wilson lines or related defects. This provides a gauge-theoretic interpretation of phases in de Sitter quantum gravity.

Detecting forest tipping points: Adapting a Landsat data pipeline for wildfire analysis in California

Prince Mutegetsi

Mentors: Christian Frankenberg and Eva-Marie Metz

California forests are increasingly affected by disturbances such as wildfire, drought, and heat. Wildfires, in particular, can prevent forests from recovering after disturbance, leading to conversion from forest to shrubland. However, it remains difficult to determine which conditions drive these conversions and which forest areas are most at risk. This project adapts a data pipeline to examine whether forests show signs of slower recovery following disturbance across California. The pipeline uses monthly vegetation observations from the Landsat mission spanning 1984 to 2025, removes seasonal patterns and long-term trends, calculates temporal autocorrelation (TAC) and variability in rolling windows, and then estimates their temporal trends using Kendall's tau. These calculations are computationally expensive, so I used Numba to parallelize the pixel-wise computations across 48 CPU

cores, enabling processing of 15.2 million pixels at 90 m resolution across the Sierra Nevada in 17 minutes. The resulting maps were compared with known forest-to-shrub conversion sites and Forest Inventory and Analysis (FIA) field plots, where the observed patterns generally matched. These results demonstrate that the pipeline can be scaled from small test regions to a California-wide analysis. Future work will strengthen validation and support development of a model that identifies forests at risk before conversion.

Sparse sensing of fluid–structure interaction dynamics using neural operators

Audrey Y. Na

Mentors: Steven Brunton, Melany L. Hunt, and Mario Sinani

Sparse-sensor reconstruction of fluid–structure interaction loads could reduce the instrumentation required for aerospace testing and enable continuous structural-load monitoring. This work investigates neural operators as virtual sensors for estimating gust-induced wing-root bending-moment histories from a small number of accelerometer measurements. The study uses an open experimental dataset containing hundreds of wind-tunnel gust encounters performed on several flexible-wing configurations under different operating and excitation conditions. Generalization is assessed using a leave-one-wing-out protocol, in which all gust-response measurements from the test structure are excluded from model training and selection. Baseline regularized linear models have been implemented to establish a benchmark for structural-load reconstruction, and neural operator architectures are developed as the primary modeling approach to assess the ability of data-driven methods to recover structural loads from sparse measurements. The project further investigates temporal encoding, causal estimation, and conditioning on structural properties to improve robustness and enable online load estimation. The causal formulation constructs a continuous load estimate using only measurements available up to the current time, enabling its interpretation as an online virtual sensor. The established evaluation framework will assess robustness across wing configurations and gust conditions, with particular attention to structural transfer, low-amplitude excitation, and generalization to unseen configurations.

Verifying a continuum approximation for architected materials in three dimensions

Yashwanth Nagarajan

Mentors: Julia R. Greer and Cyrus Fiori

Continuum mechanics theories can be used to model the mechanical response of most solid materials, including how they deform and fail. However, porous materials, such as sponges and aerogels, contain discrete structures and, as a result, cannot be directly modelled as continua. In this project, we employ a theory developed by Cyrus Fiori, my graduate mentor, that asserts that porous materials, such as architected materials, can also be modelled as continua below a critical porosity. Previous work has validated the theory in two dimensions by testing architected samples with varying porosities around this critical value in tension. However, it has not been validated in three dimensions. In order to do so, we designed three-dimensional lattices parametrically with CAD software, fabricated samples using MSLA 3D printing, and performed compression and tensile testing of the lattices to obtain global force-displacement, peak load, and critical strain value data. These data help investigate changes in the mechanical response of the material as the porosity changes, which will help validate the theory and aid engineers and scientists in the use of continuum models to analyse and design porous materials.

Detective of the skies: Unveiling the hidden population of low-nickel-mass Type IIb supernovae

Klara S. Nager

Mentors: Mansi M. Kasliwal and Christoffer Fremling

Type IIb supernovae (SNe) can display a double-peaked light curve; the first peak is generated by cooling emission after shock breakout and the second peak is powered by radioactive decay, primarily from nickel-56. As part of the Zwicky Transient Facility (ZTF) survey, we have identified a population of Type IIb SNe with very faint nickel-56-powered peaks. The low luminosities of these events may

have caused them to be systematically overlooked. This study aims to constrain the explosion physics of these faint nickel-56 peak Type IIb SNe.

We present radiation-hydrodynamic modeling of eight ZTF Type IIb SNe with very faint or undetected nickel-56 peaks. Preliminary results identify nickel-56 masses in the range $0.0059\text{--}0.025\text{ M}_{\odot}$, which is significantly lower compared to literature estimates for SNe IIb. With preliminary ejecta masses in the range $0.8\text{--}5.9\text{ M}_{\odot}$, the models produce a relatively large range of ejecta mass while maintaining a consistently low production of nickel-56. These events may offer the next steps towards reconciling the nickel-56 mass distributions of SNe II and stripped-envelope SNe, which currently lack a significant population of low nickel-56 mass SE SNe.

Investigating the nature of a serendipitous hard X-ray source in the Kepler field

Sam Nasri

Mentors: Fiona A. Harrison and Hannah Earnshaw

During observations of the Kepler Supernova Remnant, the Nuclear Spectroscopic Telescope Array (NuSTAR) serendipitously detected an unidentified hard X-ray source, Source S2. Classifying these detections is important since it helps reduce target selection bias to provide us with a true catalogue of the high-energy universe. We aim to characterize and classify Source S2 by investigating the source's spectral and temporal X-ray properties. We analysed archival NuSTAR observations of Source S2 spanning multiple epochs, performing spectral analysis to characterize its X-ray emission and timing analysis to investigate variability and periodic behaviour. Using Chandra and Swift observations, we refined the source's position and extended the analysis to lower X-ray energies. We determined that Source S2 is a faint astrophysical source detected at high X-ray energies, regularly up to $20\text{--}30\text{ keV}$, showing no significant periodic behaviour. Additionally, the source's spectrum is well described by an absorbed power-law model, consistent with non-thermal emission. The observed flux provides evidence for moderate long-term variability. The nature of Source S2 remains uncertain, with an active galactic nucleus, X-ray binary, or magnetic cataclysmic variable being the leading possibilities. Further analysis of lower-energy observations and identification of an optical counterpart will help distinguish between the possibilities.

Spectroscopic validation of candidate giant exoplanets around M-dwarf stars (GEMS)

Rachel E. Nazareth

Mentors: Dimitri P. Mawet and Andrea Lin

Giant exoplanets around M-dwarf stars (GEMS) challenge the standard core accretion model of planet formation, which predicts that low-mass protoplanetary disks will dissipate before runaway gas accretion can occur. To understand the extremes of planet formation, we need a statistically significant population of validated GEMS. However, confirming true GEMS from transit candidates identified by TESS requires spectroscopic vetting to filter out astrophysical false positives such as low-mass eclipsing binaries and brown dwarfs. We observed 24 GEMS candidates with the near-infrared WINERED spectrograph on the Magellan telescope and developed a consolidated pipeline that integrates the WINERED Automatic Reduction Pipeline (WARP), atmospheric emission-line wavelength calibration, cross-correlation function (CCF) analysis, and radial velocity (RV) extraction. The pipeline's wavelength solution and CCF performance were validated against known RV-stable reference stars. Based on CCF morphology, we categorized the candidates into single-lined and double-lined spectroscopic binaries (SB1 and SB2), as well as single-lined systems suitable for more precise RV analysis. From these classifications, we are compiling a list of false positives for public release and performing RV analysis of single-peaked targets with multi-epoch observations to distinguish planetary-mass companions from stellar/substellar ones. We have streamlined the spectroscopic vetting process for candidate surveys, enabling future population-level analyses.

GPU-accelerated radiative transfer forward model for aerosol-aware greenhouse gas retrievals

Abubakr S. Nazriev

Mentor: Christian Frankenberg

Satellite observations of CO₂ and CH₄ gases in the Earth's atmosphere are retrieved by how sunlight is absorbed as it traverses the atmosphere. Atmospheric particles (dust, smoke, sea salts, etc.) change the paths of this light through scattering. Thus, satellite retrievals of gas concentrations in the presence of such aerosols require accurate modeling of particle scattering effects, which otherwise lead to biases in retrievals. This project develops a GPU-accelerated pipeline for simulations of the high-resolution polarized spectra in three spectral regions used in observations from NASA's Orbiting Carbon Observatory (OCO) based on physical optics of scattering in a realistic particle population whose optical properties have been calculated based on the global atmospheric chemistry models. Each aerosol is defined by its size distribution, chemical composition, and altitude; its optical properties are computed according to electromagnetic (EM) scattering theory with composition-specific mixing rules and verified using independent software. The pipeline provides spectra at an instrument-independent level together with full details of the atmospheric conditions in a benchmark dataset for the purpose of training machine learning (ML) models for fast and accurate retrievals. Ongoing work is addressing the representation of particle size distributions across the discrete model categories and preparation for a year-long global generation campaign.

Utilizing full memory: An exploration of techniques in catalytic space

Tobjorn L. Nelson

Mentor: Chris M. Umans

In 2014, Buhrman et al. introduced the idea of catalytic logspace. This model differed from traditional logspace by adding an extra polysize work tape initialized with random bits that must be preserved at the end of the computation. In their seminal work, they proved catalytic logspace was stronger than traditional logspace, being able to simulate boolean circuit classes that are not known to be simulatable in logspace. Since then more striking structural results about catalytic logspace have been discovered. These include Koucký et al.'s result that catalytic logspace is equivalent to catalytic nondeterministic logspace, an open problem for its noncatalytic variant, and Alekseev et al.'s result that linear matroid intersection, a problem not believed to be contained in the circuit classes currently simulatable in catalytic logspace. We refine existing proofs of containments in catalytic logspace and work towards proving additional problems are contained in catalytic logspace.

Activating cobaltocenes for electrocatalytic CO₂ reduction to formate at low overpotential by proton-coupled electron transfer mediators

Tuyet N. Ngo

Mentors: Jonas C. Peters and David Lu

This project focuses on developing new catalytic systems for selective electrocatalytic reduction of carbon dioxide into value-added chemical products, specifically formate. In this study, a series of cobaltocene derivatives, (MenCoCp₂, n = 0, 2, 4, 6, 8, 10), serve as the homogeneous catalysts, while the piano-stool complex [CpCo(depe)H][PF₆] (Cp = cyclopentadienyl; depe = 1,2-bis(diethylphosphino)ethane) acts as a proton-coupled electron transfer (PCET) mediator. Together, these components enable the reduction of CO₂ to formate at low overpotentials with high selectivity. These cobaltocene derivatives were first synthesized and characterized and further studied through cyclic voltammetry and bulk electrolysis for formate production. Under optimized conditions, a Faradaic efficiency for formate of up to 94% was achieved at an overpotential of only 280 mV, placing this system among the best-performing organometallic catalysts in terms of both selectivity and overpotential. The mechanisms of these catalytic systems were further analyzed through UV-Vis kinetics and stoichiometric studies.

Trimodal contrastive learning for joint conditional probability estimation in fluid dynamics

Huy N. Nguyen

Mentors: Andrew M. Stuart, Claire Valva, and Hojjat Kaveh

Multimodal observations of physical systems can provide complementary information about a shared latent state, but it is not always clear whether adding modalities improves conditional inference. We study this question in Lagrangian data assimilation, using contrastive learning to align Eulerian velocity fields with Lagrangian trajectories for recovery of the underlying flow. However, extensions to more than two modalities are often constructed as sums of pairwise losses, which model conditionals such as $p(u | v)$ but need not capture higher-order conditionals such as $p(u | v, z)$. This information gap is especially significant in physical systems where individual modalities can carry marginal predictive power but offer strong information jointly.

To resolve this, we present a trimodal contrastive framework that directly quantifies joint alignment among three latent representations. We first validate our framework on Gaussian benchmarks with analytically tractable conditional distributions and mutual information, enabling direct comparison with exact higher-order distributions. We then apply the method to fluid dynamics and test whether passive tracer observations provide additional information about the Eulerian flow. Our results show that incorporating tracer data improves recovery relative to pairwise baselines, suggesting that higher-order multimodal objectives can exploit complementary physical observations not fully captured by pairwise alignment.

Faster electron density prediction with Fourier neural operator

Tuyen T. Nguyen

Mentors: Anima Anandkumar and Maurice Hanisch

Predicting the electronic structure of materials is a central task in computational chemistry and physics, but the self-consistent field (SCF) calculations underlying density functional theory (DFT) are computationally expensive at $O(N^3)$ scaling, limiting the scale and speed of simulations. Machine learning approaches, and neural operators in particular, offer a promising route to accelerating these calculations by learning to predict electron densities directly. We investigate whether an Fourier neural operator (FNO) based SCF calculation can be further accelerated by performing inference at a lower spatial resolution. We find that lower-resolution calculations can achieve similar convergence when evaluated against reference densities at the same resolution. However, when the resulting densities are compared against the original full-resolution reference, the error increases substantially. Our experiments show that most of this additional error is caused by information lost during downsampling rather than by poorer SCF convergence.

Structural analysis and reconstitution of RNA polymerase II ubiquitination

Jia-Xin Nien

Mentors: Isaac Fianu and Anjali Mirmira

RNA polymerase II (Pol II) mediated transcription is a tightly regulated cellular process that produces messenger RNA (mRNA) from DNA to control gene expression. Pol II must continuously traverse the DNA template throughout transcription, however, DNA damage creates bulky lesions that impede Pol II procession. To maintain genome stability and resume transcription, these lesions are repaired by the transcription-coupled repair (TCR) pathway. TCR removes Pol II from the lesion and facilitates DNA repair. TCR recruits the E3 ubiquitin ligase CRL4^{CSA} to polyubiquitinate Pol II. Ubiquitination can direct protein degradation by the proteasome through a three-enzyme cascade of E1 (ubiquitin-activating), E2 (conjugating), and E3 (ligating) enzymes. However, precise mechanisms behind Pol II ubiquitination and the subsequent fate of Pol II after TCR are unclear. Here, we address gaps in this understanding through establishing an in vitro Pol II ubiquitination assay. We use ligation independent cloning, bacterial and insect expression, and protein chromatography to purify ubiquitination components including an E1 (UBE1), E2 (UbcH5b), and E3 (CUL3-RBX1). This work establishes a basis for future studies of Pol II ubiquitination.

Searching for rare transients from complementary views: A multi representation normalizing-flow framework for astronomical surveys

Saina Nikmehr

Mentors: Mansi M. Kasliwal, Argyro Sasli, Ashish Mahabal, and Michael Coughlin

The Vera C. Rubin Observatory's Legacy Survey of Space and Time will transform time-domain astronomy by producing millions of transient alerts each night. Its greatest scientific promise includes the discovery of rare, rapidly evolving, and potentially unknown phenomena, yet only a small fraction of these events can receive detailed follow-up. Supervised classifiers alone cannot solve this problem: rare classes are poorly represented during training, entirely new classes have no training examples, and unknown objects may be confidently assigned to known common categories. Anomaly detection provides a way to prioritize such events by identifying light curves that do not conform to the learned population. However, the anomalies recovered by a model depend strongly on how the light curves are represented. Here, we develop a light-curve-only, multi-representation anomaly-detection framework based on normalizing flows. Separate flows learn the distributions of normal transients in two complementary spaces. The first uses physically interpretable light-curve features extracted from parametric and probabilistic reconstructions of multiband photometry. The second uses learned light-curve embeddings produced by the photometric encoder in the AppleCiDER TEMPO pipeline. Both branches operate exclusively on light-curve information, without using spectra or images. Objects with low likelihood under either normal-transient distribution receive high anomaly scores. Because extracted features and learned embeddings encode different geometries of light-curve behavior, their scores provide complementary rather than interchangeable evidence. Using archival Zwicky Transient Facility light curves as a testbed, we find that the two branches reveal different aspects of rarity. Extracted light-curve features provide physical interpretability and expose the effects of sparse sampling, observational coverage, and reconstruction quality, while learned light-curve embeddings provide stronger separation of rare events in the current experiments. The final framework is designed to combine the calibrated scores from both branches, producing a candidate ranking that is not dependent on any single light-curve representation. This approach aims to improve the scientific diversity and reliability of rare-transient prioritization under limited follow-up resources and establishes a scalable methodology for anomaly-driven discovery in the Rubin era.

Searching for recent X-ray transients in Swift/XRT data

Maya I. Nuñez

Mentor: Murray Brightman

X-ray transients (XRTs) are sources of X-rays that appear in a region of the sky where no X-ray source has been seen before. XRTs can be caused by a variety of astronomical events like Tidal Disruption Events (TDEs), which occurs when a star passes too close to a supermassive black hole. In this project, the purpose was to find and classify uncatalogued XRTs using a custom Swift archival data reduction python pipeline, NED, SIMBAD, VizieR, and Aladin. With these resources, I was able to categorize and find two uncatalogued XRTs, one of which is especially significant. LSXPS J170906.9+265122 is a newly identified Swift X-ray detection from late 2025 and appears to originate from the same Seyfert 1 active galaxy as AT2019fdr, a cosmic event surrounded by much debate. Due to its unique characterizations, it was originally categorized as a Type IIIn Superluminous Supernovae, AGN flare, or TDE. Due to my requested NuSTAR observation of LSXPS J170906.9+265122 in the 3–8 keV band, we are lead to believe that it is a late-time X-ray detection from AT2019fdr, further solidifying AT2019fdr's categorization as a TDE as opposed to Type IIIn Superluminous Supernovae (SLSN), as late-time X-ray emissions are rare for Type IIIn SLSN.

Critical mineral recovery via electrosorption across zeolite films

Cleatus R. Oakes

Mentors: Pavel A. Kots and Mark E. Davis

LTA zeolite films are investigated as selective ion-exchange and molecular-sieving materials for sustainable lithium recovery from brines. Conventional solar-evaporation extraction is increasingly viewed as environmentally and operationally unsustainable, motivating the development of more efficient separation strategies. LTA zeolites, aluminosilicate crystals with a well-defined cubic

framework and uniform pore sizes, can be synthesized as thin, continuous films on metal substrates and offer size- and hydration-based ion selectivity. In this work, LTA films were grown on nickel supports via hydrothermal crystallization. The resulting coatings were characterized for thickness, continuity, and adhesion. Preliminary ion-exchange experiments using $\text{Li}^+/\text{Mg}^{2+}$ control solutions were conducted to quantify intrinsic lithium uptake and evaluate competitive ion exclusion. The zeolite-coated electrode was then integrated into an electrochemical cell containing a $\text{Li}^+/\text{Mg}^{2+}$ control solution. Baseline transport was measured under open-circuit conditions, followed by field-driven experiments. Preliminary results show that LTA zeolite exhibits a high decationization capacity enabling strong lithium selectivity, indicating its potential as an efficient, scalable membrane material for environmentally sustainable lithium extraction. Further optimization of film thickness and applied field strength may enhance selectivity and recovery, advancing the development of more efficient lithium extraction technologies.

Supernova classification and host subtraction for DESI spectra

Esme R. Offner

Mentors: Antonella Palmese, Mansi M. Kasliwal, and Xander Hall

The Dark Energy Spectroscopic Instrument (DESI) Transients Survey captures spectra of transients, such as supernovae and tidal disruption events. We identified supernova spectra taken by DESI in the past 18 months and classified them, finding mostly Type II and Type Ia. However, supernova spectra are often contaminated by their underlying galaxy. Using DESI's vast galaxy catalog, we identified supernovae for which DESI had spectra for their host galaxies, and used these host spectra to subtract away the galactic contamination. Removing galactic contamination reveals possible circumstellar medium lines which arise when the ejecta from a supernova interacts with the mass shed by the star before going supernova. Additionally, we utilized ZTF light curves to determine the phases of the supernovae which allows us to identify infant supernovae. Using these phases, we studied how supernovae evolve spectroscopically over time by measuring the pseudo-equivalent widths (pEWs) of key spectroscopic features. The pEWs also allowed further classification of Type Ia supernovae which could lead to better Type Ia supernovae models which would make them better standard candles.

Efficient pulse-width modulation techniques for time-reversal photoacoustic tomography

Kaan Baris Oktay

Mentors: Lihong Wang and Indrasen Bhattacharya

Time-reversal photoacoustic tomography reconstructs images by recording acoustic pressure waves generated by a laser pulse within a scattering medium and re-emitting their time-reversed versions through an ultrasound array. Because each pulse provides an independent volumetric measurement, repeated acquisitions can capture dynamic biological processes as a time-resolved image sequence with reduced motion blurring. Re-emission uses a three-level switching pattern, or tri-state pulse-width modulation (PWM), rather than a smooth analog waveform, which allows the array's high-power output stage to drive the transducers strongly enough to address the signal-to-noise limitations of time-reversal playback. Each recorded signal must therefore be converted into a switching pattern that reproduces it accurately and efficiently. This project develops faster PWM-generation methods and a preprocessing pipeline that improves reconstruction quality, including contrast-to-noise ratio, peak-to-background ratio, and target visibility. A closed-form solution reproduces the existing PWM method approximately one hundred times faster. A revised preprocessing pipeline improves signal alignment, normalization, and rejection of high-amplitude ballistic components. The assumed electrical response of the array was also compared with measurements, and a compensation method was developed to restore signal strength lost as sound travels through the imaging medium. Together, these advances provide a faster and more physically accurate pipeline for higher-resolution, deeper-tissue, and dynamic photoacoustic imaging.

Phase space dynamics and systematic error analysis at the tSPECT experiment

Alexander E. Oliver

Mentors: Brad W. Filippone and Xiaozhe Zhu

Measurement of the free neutron lifetime to a high level of precision is important for multiple applications in fundamental physics, including extracting the up-down element of the Cabibbo-Kobayashi-Maskawa matrix without nuclear corrections, and the estimation of relative abundances of elements after Big Bang Nucleosynthesis. As measurements have become more precise, lifetime values taken from 'beam' and 'bottle' experiments have diverged to a statistically significant level. It is theorised that this may be due to physics beyond the Standard Model, but systematic errors in both experiments must be fully understood before any conclusions can be made. tSPECT is a 'bottle' experiment that aims to reduce systematic errors caused by neutrons colliding with walls by utilising full magnetic confinement. It is possible for neutrons with energies higher than the depth of the trap to remain confined for finite lengths of time. These 'marginally trapped' neutrons lead to underestimates of the lifetime due to the possibility of them being lost through non-decay means before the detection phase of the experiment. Hence, 'cleaning' is performed before storage to remove these neutrons. We investigate the effect of marginally trapped neutrons on the lifetime measurement, showing that with current cleaning methods the lifetime measurement will be skewed by an amount greater than the goal uncertainty of half a second. Diffuse and reflective collisional elements are simulated to demonstrate the efficacy of mixing and collimation on the cleaning phase. Significant cleaning improvements are seen when using a Compound Parabolic Concentrator to create partial collimation of the incoming neutron beam. The effect of moving geometry on the phase space is also explored, with neutron cooling on the order of 1neV observed.

Optimization of next generation lithium ion batteries with silicon graphite carbon nanotube composite as active material

Ismail B. Orak

Mentors: Sibani L. Biswal, Michael A. Vicic, Kewei Wu, and Ratnika Gupta

With the advent of ever intensifying climate change and global supply issues, a viable replacement for fossil fuels is needed for industries such as transportation. Lithium ion batteries stand as a clear contender to fill this role in technologies like electric vehicles, but currently are lacking in energy density and face long-term stability and safety issues. Silicon has long been identified as a contender for the active material of next generation battery anodes for its theoretical specific capacity of ~4,200 mAh/g, nearly ten times that of current generation graphite, but faces the debilitating issue of rapid expansion and rupturing, consuming capacity through the formation of excessive surface electrolyte interfaces. Carbon nanotubes have been introduced as both a conductive scaffold to physically restrain silicon particles from expanding and transport electrons between active materials. Ultimately, the combination of ultra fine >1 micrometer silicon and graphite particles bound in a carbon nanotube structure produces a battery with much of the capacity benefits of silicon while maintaining the longer cycle life, increased stability, and decreased costs of traditional graphite lithium ion batteries.

How to revoke users in encryption schemes anonymously

Benjamin L. Ostrovsky

Mentors: Sanjam Garg and Leonard J. Schulman

Revocation is essential for keeping encryption secure as keys leak, devices are compromised or users leave. We revisit the problem of revocation for encryption schemes and give, to our knowledge, the first approach combining anonymous, collusion-resistant revocation with unchanged keys and access rules, succinct ciphertext overhead and a revocation authority that cannot decrypt. Revoking a user publishes one random-looking pseudonym, without telling encryptors whose key was revoked. A party decrypts future ciphertexts only if it is unrevoked and satisfy the original access rule.

Our core observation is that revocation is, at heart, witness encryption for the statement that one satisfies the access rule and that the pseudonym bound to one's key has not been revoked. Although general-purpose witness encryption remains far from practical, we observe that the statements arising from revocation are simple and algebraic, and we show that they fall within the special-purpose

framework of Garg et al. (CRYPTO 2025), which compiles short arguments of knowledge into concretely efficient witness encryption using bilinear pairings. Our technical contribution is a pair of succinct non-membership gadgets for this framework, one for a single hidden pseudonym and one that certifies non-revocation jointly for all users selected by an access rule.

Both gadgets are generic and modular, as they compose with the gadgets expressing a wide range of access rules, and thus provide a single revocation mechanism that is designed once and reused across encryption schemes, without redesigning either the scheme or the revocation.

These gadgets yield the following applications. First, we give a black-box transformation adding revocation to arbitrary independently instantiated public-key encryption, with constant-size cryptographic ciphertext overhead. For unrevoked users, confidentiality continues to rely only on the security of the underlying scheme. Second, we construct revocable schemes for distributed broadcast, registered attribute-based, and registered threshold encryption, with succinct cryptographic ciphertext components and threshold decryption shares. We prove adaptive security of our schemes in the generic bilinear-group and random-oracle models.

Systematic review and pedagogical optimization of Open Stack Question Banks

Robert Otieno

Mentors: Franca Hoffmann, Jasmine R. Bryant, and Joe Champion

Open Question Banks support formative assessment in undergraduate mathematics globally. This is especially true in Kenya and Sub-Saharan Africa. One key question bank is STACK. Built on computer algebra-based marking and potential response trees, it has grown organically through many contributors leading to uneven quality and underused feedback capabilities. This project built and audit-probe-usage evidence loop rather than relying on interviews alone. It begins with a static audit checking for quality metrics per question calibrated against STACK's library. API probes distinguished genuine design limitations from analysis outputs and live usage data from deployed courses. This exposed lack of tailored branch feedback, steep retry attritions and unequal difficulty across randomized variants. These findings motivate concrete recommendations for authors, instructors and institutions. The result is a reproducible, evidence-driven pipeline for curating open question banks, with direct value to institutions and communities that rely on STACK Open Question Bank.

Removal of C55-P-P-GlcNAc-ManNAcA from MjAglH for improved structural analysis

Grace H. Otos

Mentors: William M. Clemons and Jacob Kirsh

N-glycosylation is a post-translational modification where glycans are added to asparagine residues on proteins. In humans these modifications are important for cell adhesion, migration, and signaling. In archaea they are essential for forming a protective proteinaceous layer on the membrane, enabling survival in extreme environments. The initiating proteins for this process in humans and archaea are homologs DPAGT1 and AgIH. Structural analysis of AgIH in the Clemons lab thus far has found it to be structurally similar to DPAGT1, making it of interest for comparative analysis. However, the existing structural data for AgIH shows the binding pocket to be occupied by the non-native lipid, C55-P-P-GlcNAc-ManNAcA; a result of the E.Coli expression system. This project explores techniques to obtain an apo AgIH sample to enable further study of AgIH. In order to obtain an apo sample, we applied three strategies: catalytically ablated mutants, a catalytically active mutant, and a lipid stripped wild type sample. Through mass spectrometry and lipidomics, we were able to identify the successful removal of C55-P-P-GlcNAc-ManNAcA from both catalytically ablated mutants and continue to work to purify and analyze our catalytically active and lipid-stripped samples for the removal of the lipid.

Identifying cellular targets of phage-encoded single gene lysis proteins in *Escherichia coli*

Nihira Padidam

Mentors: William M. Clemons and Roujon Nowzari

Antimicrobial resistance poses a global health threat that demands investigation of alternative target discovery. Single-gene lysis (Sgl) proteins encoded by lytic bacteriophages induce lysis in host bacterial cells. Due to their rapid evolution and diversity, the study of Sgl lysis mechanisms offers abundant opportunities to identify new druggable targets. While the number of newly discovered Sgls has increased, most Sgl targets continue to remain unknown. Here, we aim to elucidate the mechanism of lysis in *E. coli* of three newly discovered Sgls ESO 003, ESE 027, and EMM 000. We seek to develop a genetic screen to identify resistance mutations mapping to Sgl-specific targets and apply it to our Sgls of interest. Additionally, we aim to assess bacterial cellular morphology during Sgl-induced lysis to determine whether these Sgls target cell wall biosynthesis. By defining the targets of Sgl proteins, our work aims to identify bacterial vulnerabilities that can be leveraged to develop new drugs and combat the prevalence of drug-resistant pathogens.

Reconstruction of muon detector showers using machine learning

Diya Pai

Mentors: Maria Spiropulu and Soham Bhattacharya

In the CMS experiment, muon reconstruction is central to many physics analyses, and the problem grows harder with the HL-LHC upgrade, where higher luminosity will increase pileup, background rates, and detector occupancy. This project is motivated by a reconstruction bottleneck in converting low-level detector readouts into accurate particle-hit positions. The classical RecHit algorithm reaches matching efficiencies of approximately 8-25% within 1 cm and 40-65% within 5 cm depending on shower density, motivating learned methods capable of greater accuracy. We train a detector-aware transformer to predict hit positions and existence directly from detector signals. Work over the course of the project spanned diagnosing the model's weaknesses, tuning against the metrics of interest, and revising the model accordingly. Learned predictions are benchmarked against RecHit on matching efficiency, fake rate, residuals, and median and mean distance to truth, with the longer-term goal of improving reconstruction accuracy beyond the RecHit baseline.

Sample-optimal Gaussian-state tomography with single-copy nonadaptive Gaussian measurements

Bo-Syun Pan

Mentors: John P. Preskill and Senrui Chen

Quantum learning theory for bosonic systems extends finite-dimensional quantum learning to continuous-variable settings, with applications to optical interferometry, gravitational-wave detection, and proposed dark-matter searches. Recently developed tomography protocols based on collective measurements or adaptive control can achieve near-optimal performance, but require quantum memory or real-time measurement feedback. By contrast, the limits of experimentally feasible nonadaptive protocols remain largely unresolved beyond the single-mode setting.

We characterize the sample complexity of learning an unknown (n)-mode bosonic Gaussian state with energy at most (E) using single-copy, nonadaptive Gaussian measurements. We uncover a sharp dimensional transition: the optimal energy dependence is linear for one mode but becomes quadratic for every ($n > 1$), with heterodyne detection attaining the multimode optimum and adaptive Gaussian measurements achieving a quadratic improvement. For pure Gaussian states, we developed a heterodyne tomography protocol exploiting purity constraints, which achieves linear-in- (E) sample complexity and improved over mixed-state bounds. Together, these results establish sample-optimal and experimentally practical protocols for single-copy nonadaptive Gaussian tomography, while also quantifying the performance gains enabled by adaptivity.

High-throughput computational discovery of redox active transition metal mechanophores via machine learning

Joel Kai Chen Pang

Mentors: Heather J. Kulik, Jonas C. Peters, and Beck Miller

Integrating transition metal complexes into polymer networks provides a robust strategy for creating materials that transduce mechanical force into chemical responses. Mechanical stress distorts coordination geometries and perturbs redox potentials, effectively switching the reactivity of the metal center. However, navigating the vast combinatorial design space of metals and ligands limits traditional experimental discovery.

This project introduces a high-throughput computational workflow to accelerate the discovery of redox-active transition metal mechanophores, focusing specifically on salen and pincer complexes. To mitigate the challenges of experimental screening, we leverage machine learning in tandem with computationally intensive density functional theory calculations to systematically simulate molecular deformation under strain. By mapping force-modified energy landscapes, this approach predicts how physical stress drives shifts in chemical reactivity. Ultimately, using feature learning to identify key structural and electronic drivers will serve as a rational design framework for the development of next-generation force-responsive materials.

Chirality-induced effects on magnetism in helical nanowires and torus knots

Sainiket Panyam

Mentors: Charudatta Phatak, Julia R. Greer, and John Fullerton

Nanoscale magnetic helices and knotted torus structures are promising candidates for future spintronic and nanomagnetic computing devices due to their inherent chirality and curvature-induced spin textures. However, a systematic study relating geometry (pitch, radius, and knot topology) to magnetic switching behavior has not yet been established. This work uses micromagnetic simulations in MuMax3 to perform angle-dependent hysteresis measurements on helices with varying pitch, radius, and number of turns, as well as knotted torus structures under in-plane and out-of-plane fields, characterizing how geometry influences coercivity, remanence, and magnetization reversal. Decreasing helix pitch produces more gradual magnetization switching, consistent with multiple magnetic domains forming along the helix, while coercivity depends strongly on field angle, with smaller-pitch helices showing lower coercivity near perpendicular angles but a pronounced spike near 75°, indicating significant anisotropy. The knotted structures show similarly angle-dependent coercivity and remanence with a periodic pattern reflecting knot symmetry, suggesting topology-driven anisotropy could enable switching applications. In parallel, helices spanning a range of pitches and radii have been fabricated via Focused Electron Beam Ion Deposition (FEBID) and sputter deposition. Future work includes Lorentz TEM imaging of fabricated helices to compare against simulations, alongside continued simulation of knotted structures.

Evaluating the accuracy and stability of automatic differentiation of eigenvalue computations

Maria T. Papazoglu

Mentors: Michael Overton, Joel A. Tropp, and Tim Mitchell

This project investigates the use, accuracy, and efficiency of automatic differentiation to compute eigenvalue derivatives of different classes of parametric matrices. Automatic differentiation is implemented in JAX and compared with derivatives computed using classical eigenvalue perturbation theory. Experiments are performed on a variety of matrices of different sizes and properties. When using the power method to compute the dominant eigenvalue, specific matrices were generated in order to investigate any effects the power method's convergence rate might have on accuracy, stability, and computational performance of automatic differentiation. The effects of forward-mode and reverse-mode differentiation and JIT compilation on runtime and efficiency are also investigated. Different matrix parametrizations and multivariable matrix functions are also examined. These experiments are extended to Inverse Iteration and Rayleigh Quotient Iteration as well as non-symmetric matrices.

A search for Higgs boson pair production with the CMS experiment at the LHC and sensor R&D for future detectors

Nityaansh Parekh

Mentors: Maria Spiropulu and Si Xie

The Higgs self-coupling appears in Higgs boson pair production, a rare process whose measurement probes the shape of the Higgs potential and tests the Standard Model. The bbyy final state offers a favorable combination of branching fraction and mass resolution making it a leading channel for this search. But its rarity and substantial backgrounds require multivariate classifiers to separate signal from background effectively. This work focuses on validating these classifiers and using them to produce an updated statistical limit on HH production. A resampling-based method is used to check whether tightening a classifier's score cut artificially sculpts background events toward the Higgs mass peak. Validated classifier scores are then used to optimize event categorization via a multidimensional grid search, and the resulting category cuts are carried through a full statistical fit toward an updated limit on HH production. In parallel, AC-coupled Low Gain Avalanche Diode (LGAD) timing sensors were characterized for a separate detector R&D effort, studying charge sharing between neighboring pixels and time resolution as a function of bias voltage.

Invariant measure of the half-space polynuclear growth model

Jin Park

Mentor: Xincheng Zhang

Polynuclear growth (PNG) is an integrable probability model that describes random surface growth behavior in space-time. Its state space is given by the set of upper semi-continuous height functions from $R \rightarrow Z \cup \{-\infty\}$. The dynamics consist of nucleation events governed by a space-time Poisson point process with rate 2 and deterministic growth of up and down jumps in the height function with rate 1. The invariant measures for full-space PNG are constructed by the superposition of two independent Poisson point processes with intensities ρ and ρ^{-1} . However, the invariant measures for the half-space model, with height functions on $R_{\geq 0}$, have not been identified. We prove that μ_ρ from Matetski, Quastel, and Remenik with $\rho = \lambda^{-1}$, where λ is the rate of nucleation at the origin, is invariant for half-space PNG.

Quantifying Gaussian decompositions of subgaussian random vectors and towards the convexity conjecture for solid sets

Caleb Pascale

Mentor: Antoine Song

We explicitly quantify the value of the uniform rescaling $\kappa \in \mathbb{R}$ present in the positive answer to the subgaussian vector problem, which states that for any centered 1-subgaussian random vector X in $\mathbb{R}^n$, we may write $\kappa X = G_1 + G_2 + G_3$, where $G_i \sim \mathcal{N}(0, I_n)$.

This result on sums of Gaussian vectors is equivalent to an affirmative answer to Talagrand's convexity problem. We then employ both geometric and probabilistic methods to obtain intermediate results towards a version of the conjecture for solid sets in $\mathbb{R}^n$.

Synthetic data-driven knowledge integration for large language models

Amrita Pasupathy

Mentors: Yisong Yue and Lauren Hyoseo Yoon

Although large language models (LLMs) are powerful tools for accessing information, they often struggle to incorporate new knowledge acquired after pretraining without expensive retraining. To address this issue, we propose a knowledge integration framework that can add new information to a base model while also preserving the ability to incorporate future information incrementally. Specifically, we construct a synthetic data engine that transforms the original data into synthetic variants designed to better expose hidden insights and relationships. We then train a parametric

memory module to encode this knowledge in its parameters. Finally, we evaluate our approach by examining the balance between memorization and generalization, with the goal of developing a scalable system that effectively integrates new knowledge while maintaining strong performance on downstream tasks.

Fusing sequence identity and protein language model embeddings with machine learning for fungal pathogenicity phenotype prediction

Ridham V. Patel

Mentors: Ashish Mahabal, Nitin Kumar Singh, and Swati Bijlani

The fungal kingdom is known to harbour a significant portion of dark matter, where the phenotypes of fungi are unknown. Traditional methods like BLAST and HMMER use sequences to capture similarity between classes representing protein and organism genes, which are not sufficient for making claims about phenotypes. On the other hand, modern protein-language foundation models consider structure and capture a higher form of functional similarity, while smoothing out small differences that traditional methods can catch. Therefore, making more confident claims about fungal phenotypes requires bringing about the best of both worlds, which we are addressing in our work using machine learning. We develop a label-free, self-supervised machine learning technique, GMMLOSO, that can weigh all available information sources and provide a provably confident combined orthology score. Additionally, we increase our confidence in the prediction by stacking it with a supervised learner. Evaluated across 158 reference proteins and approximately 1,525 organisms spanning 17 phenotype categories, the stacked model achieves 0.961 AUROC compared to 0.626 for BLAST alone. These computational results warrant further biological validation to confirm the predicted phenotype assignments.

Design and evaluation of architected heat exchangers

Akanksh Vishwan Patil

Mentors: Angkur Shaikeea and Saksham Sharma

The exponential rise of computing infrastructure has drastically increased heat generation, presenting a huge challenge for cooling systems. Modern heat exchangers must operate under strict volume and weight constraints while maximizing heat extraction efficiency and minimizing base temperatures. This project looks at the viability of using architected materials, materials that derive their properties from their structure or 'architecture', to satisfy these demands. Utilizing an entropy generation minimization framework, the thermodynamic efficiency of various heat exchanger topologies is analyzed. We compare these distinct topologies across a range of relative densities to determine if architected structures offer superior thermal performance over conventional designs. Finally, we provide design principles which engineers could use for making electronic cooling systems using architected materials.

Change of variables and difficulty measures for stable Andrews–Curtis search

Avigya Paudel

Mentors: Sergei G. Gukov, Lucas Fagan, and Michele Tarquini

The Andrews–Curtis conjecture, open since 1965, asks whether every balanced presentation of the trivial group can be simplified to a trivial presentation by a fixed set of moves; hard cases such as $AK(3)$ have resisted every attempt. We study the conjecture's stable version via a change of variables (CoV) that stabilizes and destabilizes with arbitrary words (Lemma 11): a single "supermove" whose exact AC-move count is unknown, making it powerful enough to potentially solve previously unsolved presentations. On a benchmark of 66 presentations, the best transformed start solved 52 instead of 40; on the 40 both solved, it cut the search effort more than sevenfold and roughly halved the solution length. Automorphisms and AC moves reduce the 261 previously unsolved presentations to 124 equivalence classes; however, none were solved even with exhaustive CoV search. A random transform performs no better on average than direct search; the gain comes from generating many

starts and keeping the best. A simple block-repetition measure separates solved from unsolved cases with 94 percent accuracy; a related tangling measure remains low for solved cases; and steering the search with these measures raises solves from 29 to 38 of 60. This offers an opportunity for machine learning: selecting the right change of variables at the right presentation state.

Reducing blade-wake interaction noise from upstream propeller shroud geometries on M4 Phoenix

Anshi Paul

Mentors: Mory Gharib, Matthew Anderson, and Reza M. Nemovi

The Multi-Modal Mobility Morphobot (M4) Phoenix is a morphing robot that combines aerial and ground locomotion using propellers surrounded by treads. Its high acoustic signature, however, may limit operation near people. This project tested whether redesigning the structural shrouding upstream of the propeller could reduce noise while preserving the geometry required for ground mobility. Interchangeable 3D-printed ABS shrouds varied spoke count, straight versus 45° Archimedean spoke paths, and the addition of a rounded leading edge and tapered trailing edge. A single-propeller rig recorded optical RPM, A-weighted sound level, and microphone spectra across five operating points, and knock tests identified structural vibration modes. At the highest operating point, the no-shroud baseline averaged 107.0 dBA and straight-spoke shrouds averaged 119.5 dBA, whereas the lowest measured shroud configuration, the rounded ten-spoke Archimedean design, averaged 115.4 dBA, a 4.1 dBA drop that closed 32.8% of the measured dBA gap between the current straight shroud and the unshrouded baseline (61% reduction in the linearized A-weighted sound energy). Spoke count produced little measurable change, while curving the spokes reduced overall sound pressure level and weakened several blade-passing-frequency harmonic peaks. These results identify dephasing blade-wake interactions as a promising strategy for reducing M4 Phoenix noise.

Sharpening the search: Optimizing missed NEO detections for NEO Surveyor

Nikhita Penugonda

Mentors: Joseph Masiero and Ishan Mishra

NASA's upcoming mission NEO Surveyor aims to significantly expand the catalog of Near-Earth Objects (NEOs) by detecting fainter bodies using observations in two bandpasses, NC1 and NC2. A key challenge in the completion of the software is optimizing the NEO Surveyor Survey Data System (NSDS), with a specific focus on its Moving Object Detection Pipeline (MODP). MODP links individual detections into triplets that encompass the object's movement throughout the observation period, known as a quint. This project investigates why simulated NEOs are consistently missed by NMODE, the NEOS Moving Object Detection Engine, and identifies opportunities to improve efficiency of the pipeline. Using the simulated survey data, diagnostic scripts were developed to analyze the characteristics of missed detections, including average velocity, trajectory curvature, source crowding, moving point spread function accuracy, and more. Preliminary results indicate that the missed objects tend to cluster around an average velocity of 1deg/day and exhibit larger trajectory RMS values, suggesting that the current velocity grids and triplet allowance thresholds may need to be tweaked. These findings will guide refinements to pipeline parameters, with the goal of improving MODP completeness and supporting the mission objective of reliably detecting at least 90% of simulated NEOs.

Fast configuration-space convex set computation for optimal motion planning

Deon F. Petrizzo

Mentors: Aaron D. Ames, Blake Werner, and Taoran Liu

In optimization-based motion planning, one of the major problems is dealing with challenging nonconvex collision-avoidance constraints. In this work, we present an algorithm that solves this problem by efficiently generating collision-free convex sets in robot configuration space via parallel computation. Given a collision-free piecewise-linear path from a fast sampling-based planner, the algorithm inflates obstacle-free convex regions around points sampled along the path and links overlapping regions into a graph. A shortest-path search fixes the sequence of regions into a convex corridor, reducing trajectory optimization to a single convex program whose solutions are certified to

remain collision-free. We then instantiate this algorithm in a global multi-modal planning pipeline, in which free space depends on the locomotion mode of the robots. Aerial convex corridors are inflated from raw lidar point clouds, and ground corridors are generated from traversability-derived occupancy maps, both using the same general corridor inflation algorithm. We then optimize the trajectory through these collision-free sets. The efficacy of our approach is demonstrated through the successful deployment of this pipeline in a hardware demonstration involving both a humanoid and the Multi-Modal Mobility Morphobot, as well as through the establishment of conditions under which the corridor inflation algorithm is probabilistically complete.

Wave attenuation and hydrodynamic analysis of artificial rocky reefs

Kyra T. Phaychanpheng

Mentors: Christine M. Baker and John O. Dabiri

The accelerating pace of coastal environmental change demands innovative solutions to protect vulnerable shorelines against erosion. As sea levels rise and surface gravity waves intensify, wave attenuation—the gradual dissipation of wave energy within the wave boundary layer, enhanced by seabed roughness—becomes a critical defense mechanism. Artificial reefs leverage this and provide substrate for outplanted kelp, yet their hydrodynamic effects remain poorly characterized. To address this gap, we measured wave attenuation across three experimental conditions in a wave-current flume: a smooth reef, a rocky reef without kelp, and a rocky reef with a model *Egregia menziesii* canopy. Spectral analysis was performed on wave-gauge data to determine transmission coefficients and energy levels for each condition. Measurable differences were observed, though statistical significance was limited due to flume conditions. To investigate the dissipative flow structure within the boundary layer, we utilized a particle image velocimetry (PIV) laser system to capture full velocity fields above the reef. This enabled calculation of boundary layer thickness, Reynolds shear stress, and turbulent kinetic energy, including spatial maps of coherent wake structures generated by the reef roughness. These results provide new insight into the unexplored physical mechanisms by which rocky reefs can help shores dissipate high wave energy.

Creating benchmarks for robust space robotics decision making with reinforcement learning

Ethan Pichon

Mentors: Benjamin Riviere and Gunter Niemeyer

Space exploration increasingly depends on on-orbit servicing (OOS), since human extravehicular activity is risky, costly, and often infeasible. Autonomous OOS manipulation lacks standard tasks capturing its defining difficulty, a free-floating base whose reaction couples into every contact. We present three contact-rich benchmarks, satellite refueling, uncooperative stabilization, and wire management, specified in simulation and fabricated for a microgravity mobile manipulator testbed. Baselines trained with PPO under hybrid position and impedance control insert a probe with 4.3 mm clearance per side on 23 of 24 episodes. We further find that transfer is set by the benchmark specification itself: observing only measurable quantities and acting in the vehicle's native command space reproduces simulation exactly, while an otherwise identical policy trained on privileged state loses half its performance. These benchmarks establish baselines for an online RL planner that uses state-space mixing to explore bottleneck states.

Sub-grid black holes: Learning the multi-scale dynamics of accretion feedback

Yash Pincha

Mentors: Anima Anandkumar, Chuwei Wang, Elias R. Most, and Haiyang Wang

Supermassive black holes co-evolve with the galaxies that host them through a two-way feeding-and-feedback loop, but the physics involved spans nine orders of spatiotemporal magnitude. Near the event horizon, turbulent magnetized gas accretes and launches relativistic jets and winds; far out, at galactic and intergalactic distances, those same outflows deposit energy that regulates how gas cools and stars form. Simulating the whole system at once with a single high-fidelity magnetohydrodynamic model is computationally infeasible, since resolving the black hole forces an extremely small timestep everywhere. Existing workarounds instead inject a fixed, hand-tuned feedback recipe near the black hole, but such static recipes are not dynamically consistent with the surrounding flow and cannot

reproduce the flickering, time-variable behavior of real feedback. We develop a machine-learning alternative: two neural operators that decouple time evolution from spatial structure. One operator evolves the flow state on a boundary sphere forward in time, trained on short simulated trajectories yet able to run stably for tens of thousands of steps beyond that training window. A second operator, the focus of this work, reconstructs the full three-dimensional interior flow from a single boundary snapshot, conditioned on global accretion diagnostics. Alternating between the two forms a "cyclic zoom" that couples small and large scales without the mismatch of fixed recipes.

Understanding the mechanisms behind sex-biased trophectoderm advantage in preimplantation development using stem cell-based embryo models

Manasvi Pinnaka

Mentors: Magdalena D. Zernicka-Goetz and Sergi Junyent

In vitro fertilization (IVF) is commonly associated with a higher male-to-female birth ratio compared to natural conception, but the developmental mechanisms underlying this sex bias remain poorly understood. During early development, both mouse and human embryos form the blastocyst, a structure comprising three lineages: epiblast (EPI), primitive endoderm (PrE), and trophectoderm (TE). The TE, which gives rise to the placenta, has emerged as a key focus of how sex-specific differences arise. Meglicki et al. (in revision) demonstrated that male mouse embryos exhibit accelerated TE development relative to females, specifically in high glucose culture, from the morula stage onwards. To further investigate the relationship between sex and TE function, this study utilized blastoids, stem cell-based embryo models that recapitulate key features of the natural blastocyst. Male and female TSC lines were used to model the TE and combined with pre-established male and female ESC lines, representing the EPI, and inducible Gata4-expressing cells, representing the PrE. Blastoids with different combinations of male and female lineages were generated to determine how the sex of each lineage influences early developmental fitness. Blastoid formation and cavitation efficiency as well as the composition and organization of the resulting cell lineages were compared across sex combinations.

Engineering fluorescent proteins for robust magnetic field sensing in mammalian cells

Francesco Plastina

Mentors: Mikhail G. Shapiro and Hengyu Li

Abstract withheld from publication at mentor's request.

Scaling full-3D inverse design to large-area photonic systems

Alexander P. Plekhanov

Mentors: Kiyoul Yang and Kerry J. Vahala

Inverse design uses gradient-based optimization of Maxwell's equations to discover compact nanophotonic device geometries with prescribed optical responses, inventing structures beyond conventional human design intuition. However, the computational cost of full-wave 3D simulation has generally limited designs to footprints below roughly 30x30 μm . This restricts their performance and their use in complex, highly connected photonic systems.

We developed a fabrication-aware inverse-design framework that substantially expands the accessible design area to approximately 110x110 μm on a single accelerator and 200x200 μm using distributed computation at telecommunications wavelengths. This advance enables large multiport photonic systems, rather than isolated components, to be optimized directly.

We apply the framework to SiN and SOI devices, including a wavelength-division multiplexer and a broadband 50:50 splitter. A central target is a nanophotonic neural-network accelerator implementing a 36x36 optical matrix transformation in a single structure. By performing the core linear operations of neural-network inference directly in light, such devices could provide massively parallel, high-throughput computation with reduced data movement and energy consumption, while substantially exceeding the scale of previous inverse-designed optical matrix processors.

The role of slow slip in injection-induced seismicity: Application to the Soultz-sous-Forêts geothermal site

Kyrill Pluquet

Mentors: Jean-Philippe Avouac, Alexis Saez Uribe, and Samson Marty

Injection-induced seismicity refers to earthquakes triggered by fluid injection into the subsurface, particularly during geothermal operations. This study focuses on the role of slow fault slip, which can redistribute stress beyond the region directly affected by fluid pressure. A numerical model was developed by coupling a slow-slip model with a recent formulation of the Dieterich seismicity-rate model. Using injection flow rate or well-pressure data as input, the model calculates pressure diffusion along the fault, the resulting slow slip and stress changes, and their effects on earthquake occurrence rates. The model was applied to the 1993 hydraulic stimulation experiment at Soultz-sous-Forêts, France. Its predictions were compared with the observed spatial and temporal distribution of earthquakes and the cumulative number of events. The initial model reproduced the spatial envelope of seismicity and the cumulative event count but did not explain a low-activity zone near the injection point. A simplified representation of earthquake nucleation length was therefore introduced. The updated model shows good agreement with the spatial distribution, cumulative number, and temporal evolution of the observed events, despite its simplified fault geometry. Future research will investigate poroelastic effects, improve pressure modeling and fault geometry, and assess the model's ability to forecast induced seismicity.

Dynamic control of solidification conditions via a movable platform for scalable freeze casting

Benjamin D. Porosoff

Mentors: Katherine T. Faber and Wesley D. Patel

Freeze-casting is a materials processing technique that leverages the directional solidification of solutions and suspensions to form porous materials through sacrificial templating. Controlling the temperature gradient and the freezing-front velocity during the freezing process dictates the finished materials properties' and pore morphology. This project focuses on the development of a continuous movable freeze-casting setup which allows for much longer samples compared to previous static freeze-casting setups, developing the process toward production scale. This apparatus physically decouples the temperature gradient and freezing-front velocity, enabling the exploration of regimes that are inaccessible to current freeze-casting apparatuses, and allowing for the creation of novel crystal morphologies and freeze-cast materials. The first part of this project explores the apparatus design, and testing of the control system and validation against current materials and parameters. The final part of this project explores previously inaccessible regimes of the parameter space. Lastly, production scale-up to larger samples is demonstrated with the movable technique which uses a constant rate of energy to form samples, regardless of their length. Additionally, design of materials with position-dependent morphology is demonstrated via the physical decoupling which allows the machine to form continuously frozen samples while altering their morphology with length.

Magnitude of metric spaces and spectral action models of gravity

Peter B. Powell

Mentor: Matilde Marcolli

The spectral action and the magnitude are numerical invariants with very different origins: the first is the trace $\text{Tr}(f(D/\Lambda))$ of a cutoff applied to the Dirac operator of a spectral triple, while the second is a category-theoretic generalization of the Euler characteristic to metric spaces. For a homogeneous Riemannian manifold both have similar asymptotic expansions in which the coefficients are integrals of certain curvature forms over the manifold, suggesting a relationship not apparent from the definitions. For finite metric spaces, we construct a spectral triple recovering the metric through Connes' distance formula and show that both invariants can be written as sums over paths. For spheres we compute both exactly and expand each by Poisson summation with remainder estimates. The magnitude of the sphere of radius R is a polynomial in R up to exponentially small corrections, and the multiplicities of the positive Dirac eigenvalues on a sphere are the values of a polynomial evaluated at these

eigenvalues. We show that the magnitude polynomial for the n -sphere and the multiplicity polynomial for the $(n+1)$ -sphere are proportional after rescaling R to $iR/2$. This gives a direct relation between coefficients of the magnitude and spectral action expansions for spheres. We also investigate whether a more general relationship exists for all homogeneous Riemannian manifolds.

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.

Bacteria growth and chemotaxis in the presence of fixed heterogeneous nutrient sources

Anna Prenowitz

Mentors: Sujit Datta and Tanumoy Dhar

Groups of actively moving particles, ranging from swimming microorganisms to synthetic self-propelled colloids, often exhibit collective chemotaxis — directed motion along a chemical gradient that the constituent particles generate themselves. This self-generated gradient enables active collectives to migrate as coherent fronts that travel at constant speed over large distances. Such behavior is well characterized in homogeneous environments, and continuum frameworks coupling bacterial density to nutrient consumption are well studied in porous media. In natural environments such as soil, however, nutrient gradients are not set solely by local consumption; they may be imposed externally by heterogeneous sources, including plant roots and variable chemical composition. In this work, we theoretically study collective chemotaxis in the presence of externally imposed sources of chemoattractants and chemorepellents, modeled as dynamic fields consumed by bacteria. We model the system and solve the governing equations by the finite element method in FEniCSx, and use the framework to map the parameter regime. We determine how a localized chemoattractant source governs bacterial dynamics and how imposed sources can steer bacterial organization. Our work thus establishes a route to controlling bacterial chemotaxis in complex environments.

Bandwidth-aware collaborative mapping between a humanoid and an aerial robot

Raymond A. Provost

Mentors: Aaron D. Ames, Erina Yamaguchi, and Ryan M. Bena

This project develops a real-time communication and shared-mapping system that enables an aerial robot to extend the perception and planning horizon of a humanoid robot without continuously transmitting high-volume sensor data. Each robot uses 3D LiDAR measurements to construct an occupancy map within a shared reference frame. Rather than transmitting complete point clouds or maps, the aerial robot transmits compressed updates representing changes to occupied and cleared regions over a low-latency wireless link. A second, reliable communication link periodically synchronizes the robots' maps to correct errors accumulated from lost or incomplete updates. This hybrid architecture combines the longer range and reliability of a lower-frequency link with the higher bandwidth of a higher-frequency link. Both robots localize within the shared map before their observations are integrated, allowing spatial information collected from different viewpoints to be consistently combined. The communication and mapping architecture was implemented and evaluated using physical hardware, LiDAR sensors, and wireless radios as a precursor to deployment on the target robotic platforms. By reducing the amount of information required to maintain a shared representation of the environment, this architecture extends the spatial context available for motion planning and may aid in coordinating behavior among heterogeneous robotic systems.

Constraining thermal evolution models of Europa using a Markov Chain Monte Carlo framework

Nat C. Pujet

Mentors: Jonathan Lunine, Flavio Petricca, and Kevin Trinh

The properties and potential habitability of Europa's subsurface ocean depend on the moon's thermal history and resulting degree of differentiation. Modeling the thermal evolution of Europa's interior over 4.5 billion years lets us test which pathways produce interior structures matching observations, placing new constraints on the moon's evolution. Previous studies have identified thermal evolution scenarios that match Europa's derived moment of inertia (MoI) and observed radius. However, the results of thermal models depend on many uncertain parameters, including formation time, composition, and tidal heating, making it taxing to search the large parameter space for the most consistent thermal histories. In this work, we use an MCMC framework to efficiently explore many 1-D thermal evolution models of Europa and constrain them using Galileo measurements of MoI and radius, yielding probabilistic insights into parameters that govern Europa's thermal history. The resulting posterior distributions reveal a degeneracy among initial metal fraction, rock hydration, and mantle density, indicating that multiple compositions can reproduce Europa's bulk properties, while formation time and sulfur content remain largely unconstrained in the current MCMC implementation. This work establishes the viability of combining thermal models with Bayesian parameter estimation, an approach that could be generalized to other icy satellites.


Approaching the total synthesis of Phorbasone A

William R. Putzer, Jr.

Mentors: Brian M. Stoltz and Sean Corman

In this paper, we aim to conduct a total synthesis of Phorbasone A, a complex natural product isolated from marine sponges. We explore a synthetic pathway that is currently proposed to start with cheap, commercially-available materials (R)-carvone and dimethylcyclohexanone. A series of chemical reactions transform the starting materials until the resultant fragments are at last joined together via

a Suzuki coupling. Synthesis of the western fragment (originating from dimethylcyclohexanone) has been demonstrated, while synthesis of the eastern fragment (originating from (R)-carvone) is still under speculation and experimentation.



Transition from inspiral to plunge for eccentric EMRIs with spinning secondary

Oleksandra Pyshna

Mentors: Katerina Chatziioannou, Lisa V. Drummond, and Lucy M. Thomas

Black hole binaries with extreme ($\sim 10^5$) mass ratios that inspiral due to the gravitational-wave (GW) backreaction are known as extreme mass-ratio inspirals (EMRIs). Such signals are important targets of the upcoming LISA (Laser Interferometer Space Antenna) mission, as they enable precision tests of strong-field gravity and detailed measurements of black hole properties. At the inspiral stage, the secondary's motion can be treated as a trajectory through a sequence of geodesics, with GW backreaction slowly evolving its orbit. As the secondary approaches LSO (Last Stable Orbit), GW backreaction becomes stronger and overcomes the geodesic orbital restoring force, leading to the plunge of the smaller component. In this work, we obtain the first complete inspiral-transition-plunge worldlines for eccentric EMRIs with the spinning secondary by employing the novel approach of the 'virtual worldline' parametrization of the spinning secondary trajectory. Furthermore, we carry out an analysis of how the spinning secondary alters the transition regime and identify at which mass ratio this effect becomes negligible. From the derived worldlines, we model GW waveforms where we quantify the spinning secondary effect on the observable GW emission and assess its expected SNR (Signal-to-Noise Ratio) for the future GW detector LISA.

Aerosol corrected methane retrievals from slant viewing geometries

Ellin J. Qiao

Mentors: Christian Frankenberg and Jiani Yang

Methane is a greenhouse gas that directly motivates global warming and climate change. Satellite and ground-based measurements of Earth's atmosphere have been leveraged to retrieve top-down measurements of column-averaged dry-air mole fractions of greenhouse gases. The dominant source of error for these observations is uncertainty due to light path deviations caused by aerosol scattering. One method to diminish aerosol-induced errors is to use a reference gas with a well-known dry mixing ratio as a light-path proxy to cancel out scattering effects. In this study, we investigate the effects of aerosols on greenhouse gas retrievals from slant viewing angles. A non-scattering retrieval is used to quantify errors arising from neglecting aerosols. We then determine the geometries and aerosol loadings where the N_2O light-path proxy is suitable and assess when full-physics retrievals are necessary due to non-linear aerosol effects. We report that nonlinearities due to variations in absorption strength dominate over discrepancies in aerosol effects at different wavelengths for shorter slant paths, but find that the N_2O proxy is more effective for very slant viewing geometries. We also apply the N_2O proxy to data from the California Laboratory for Atmospheric Remote Sensing-Fourier Transform Spectrometer (CLARS-FTS) and compare it against current retrievals.

Combustion in porous media burners: Computational and experimental investigation of thermo-structural material response to predict burner fracture locations

Hannah V. Ramsperger

Mentors: Matthias Ihme and Michael Mello

Porous media burners are a promising technology to enable stable combustion under fuel-lean conditions in gas turbine engines, reducing nitrogen oxide and carbon monoxide emissions while improving fuel efficiency. However, experimental observations with silicon carbide foam materials show that these porous ceramics experience severe thermo-mechanical loading during combustion, leading to brittle fractures, thereby limiting durability. Therefore, a computational framework predicting thermo-mechanical stresses and potential failure locations was developed using thermo-structural finite element analysis in ANSYS Mechanical. Given silicon carbide's high compressive and relatively low tensile strength, failure is evaluated using the Rankine criterion, which assumes fracture occurs when the maximum principal stress exceeds the ultimate tensile strength. Considering an engineered matrix topology of a triply periodic minimal structure, the simulations predict pyramidal planes of critical stresses under an applied temperature gradient obtained from a porous media combustion heat transfer simulation, with the highest stresses occurring at the node-to-strut transition. These stress patterns are qualitatively consistent with experimentally observed pyramidal fracture surfaces, where fractures propagate perpendicular to the strut axis at the strut-node junction. These results suggest that the model captures key thermo-mechanical loading mechanisms and provides a framework for studying geometry and topology optimization of burner designs to improve durability.

Spontaneous chiral symmetry breaking of counter-propagating modes via the Kerr nonlinearity within ultra-high-quality factor optical resonators

Avani Ranka

Mentors: Kerry J. Vahala and Haojing Chen

Optical gyroscopes are essential to modern navigation, offering precision and stability unmatched by their mechanical counterparts. Chip-based ring laser gyroscopes are also compact and lightweight with a low lasing threshold. However, at low rotation rates, counter-propagating modes couple via backscattering and lock to a single frequency, making small rotations impossible to measure. We present progress toward a symmetry-based method for breaking this lock-in effect without the external hardware usually required. Our approach uses the Kerr nonlinearity to break chiral symmetry between counter-propagating waves within a silica-based optical resonator, which splits their resonance frequencies. We bidirectionally pumped an ultra-high-quality factor resonator with 1550 nm laser light to form a symmetric standing-wave state. We then swept the laser frequency across modes with quality factors exceeding 100 million at increasing input powers while monitoring the transmitted power in both directions. Current measurements of resonators whose mode spacing precisely matches the Brillouin shift demonstrate counter-propagating modes in the symmetric state even above the predicted threshold, because Brillouin lasing clamps the circulating power. We now intend to use devices with mode spacings detuned from the Brillouin shift and increase the power coupled into the mode until the imbalance between counter-propagating powers is clearly resolvable.

Simulation-trained neural reconstruction of ultrasound data for detecting brain micrometastases

Tanish Rao

Mentors: Anima Anandkumar and Ryan Han

Metastatic tumors in the brain are usually detected only after they have grown well beyond the size at which treatment is most effective. Ultrasound imaging with gas vesicles, gas-filled protein nanostructures that scatter sound and can be expressed by tumor cells, offers a route to earlier detection, but skull-induced aberration degrades conventional beamforming of small, sparse targets. We are developing a learned reconstruction that maps raw radiofrequency data from cross-amplitude modulation acquisitions directly to gas vesicle concentration and location, bypassing beamforming entirely. Because training data of this kind cannot be acquired experimentally at scale, we generate it by full-wave acoustic simulation: coronal mouse skull slices from micro-computed tomography are

fitted with a polymethylpentene acoustic window matched to the experimental setup, converted to speed-of-sound maps with a delineated brain cavity, and seeded with gas vesicle inclusions from which the returning radiofrequency signal is computed. A Fourier neural operator trained on these simulations recovers the position and concentration of inclusions across the full skull, including cluster configurations not represented in training. Ongoing work extends the simulation to heterogeneous cluster morphologies and validates the model against experimentally acquired radiofrequency data.

Identifying underlying glioblastoma treatment resistance using computational modeling of spatially resolved CRISPRi perturbations at a single-cell resolution

Krish U. Raval

Mentor: Matthew W. Thomson

Glioblastoma is one of the most aggressive and lethal forms of cancer, characterized by its rapid proliferation, and resistance to conventional therapies. The standing complexity lies within the diverse cellular heterogeneity within tumors as well as its distinct number of variable gene programs that can resist effective treatment. As the tumor microenvironment provides a collection of intricate pathways and barriers that prevent regular studies to understand tumor growth, recent advancements in spatial-omic technologies have been made to gather insight about the spatial context of cancer. Leveraging Spatial Perturb-seq on the platform OpenST, our project aims at computationally defining cell state transitions across spatial niches in glioblastoma under CRISPRi perturbations. We analyze primary and recurrent human glioblastoma samples that have been profiled using spatial transcriptomics at a subcellular resolution while also maintaining transcriptome wide profiling. Through this approach we characterize and evaluate potential therapeutic hits that can be used in treating GBM resistance at a single-cell level, through statistically benchmarking perturbations in established microenvironments within GBM.

The gaussian curvature of minimal surfaces in S^n

Tisya C. Rawat

Mentor: Riccardo Caniato

We study the problem of constructing closed minimal surfaces with strictly negative Gaussian curvature in the round sphere S^4 . Such surfaces are now known to exist in spheres of sufficiently high dimension, but their existence in S^4 remains open. Our previous work established an obstruction within the superminimal setting: every smooth superminimal immersion of a closed surface of genus $g \geq 2$ into S^4 must have an umbilic point and therefore cannot have strictly negative Gaussian curvature everywhere. During this project, we developed an alternative proof of this result and investigated possible constructions outside the superminimal framework.

Our current approach seeks large-dimensional families of equivariant minimal immersions. For a closed surface Σ_g , we consider representations $\rho : \pi_1(\Sigma_g) \rightarrow O(5)$ belonging to the smooth Zariski-dense locus of the corresponding character variety, which has dimension $-\chi(\Sigma_g) \dim O(5) = 20(g - 1)$. We aim to construct ρ -equivariant branched minimal immersions $\Sigma'_g \rightarrow S^4$ that vary smoothly with ρ . One possible strategy to achieve our goal is to fix a conformal structure on Σ_g , minimize the equivariant energy, and then vary the conformal structure in search of a conformal harmonic map. This method becomes increasingly difficult as the genus grows: the Teichmüller space of Σ_g has dimension $6g - 6$, so a growing number of conformal directions must be controlled simultaneously, while its noncompactness allows conformal structures to degenerate by developing long, thin regions. We therefore investigate whether equivariant optimization of the first Laplace–Beltrami eigenvalue offers a more effective alternative. Related eigenvalue-optimization methods have successfully produced embedded minimal surfaces in S^3 across arbitrary genera, suggesting that they may better accommodate the increasing complexity of the domain. The main challenge is to ensure that the resulting minimal immersions have neither branch points nor umbilic points.

Optimization of patterned gold-assisted exfoliation for monolayer transition metal dichalcogenides

Alexander M. Recce

Mentors: James Hone, Keith C. Schwab, and Simon Yoon

Mechanical exfoliation yields high quality, research-grade 2D transition metal dichalcogenide (TMD) crystals; however, Scotch-tape exfoliation produces flakes of unpredictable size and position. We demonstrate a gold-assisted exfoliation approach that leverages Au-TMD adhesion to yield large-area, uniform monolayers. Using a patterned lithographic template with exposed SiO₂ regions, the hybrid gold substrate exfoliates flakes with high yield and predetermined geometry. To optimize the template geometry, we identify geometries that minimize mechanical deformation by assessing wrinkle formation by shape with a U-Net-based density-estimation pipeline. MoS₂ flakes are exfoliated on the hybrid gold substrate and stacked with polycarbonate-based dry pickup to produce zero-angle, artificial 3R-MoS₂ samples with thicknesses of 3 layers (L), 5L, and 10L. Measuring polarized second-harmonic generation (SHG) in regions of 1L to 5L indicated six-lobed angular responses and aligned crystallographic orientations. SHG intensity increased with layer count, with 41.8 times the monolayer response at 5L. A and B exciton resonances appeared in energy-dependent SHG spectra, but no interlayer exciton peak was found. However, there appeared some increased SHG intensity near the interlayer exciton energy range. Together, hybrid-gold tape facilitates reliable stacking of TMD monolayers at desired twist angles, enabling scalable assembly of 2D systems with modifiable structure and electronic properties.

Energy and momentum transport through 3D-printed graded auxetic lattice structures subjected to impact

Carlo A. Repetto

Mentors: Guruswami Ravichandran and Aditya B. Shedge

Lattice structures have been widely studied under both quasi-static and dynamic loading due to their superior mechanical properties and high energy absorption. However, the effects of geometric gradients throughout the structures are much less understood, especially for the mechanical response of auxetic lattice structures under compression. In this study, we investigate the effects of gradients in ligament thickness and ring radius on the behavior of hexachiral lattices (lattices with a negative Poisson's ratio) under quasi-static and dynamic loading. We use an automated workflow that generates graded hexachiral lattices, performs finite element analyses (FEA), and trains a machine learning (ML) model to predict the ideal geometric gradients for peak transmitted stress, specific energy absorption, and strain at peak stress. We then fabricate the optimal geometries using resin 3-D printing and experimentally test the samples with an Instron machine and a gas gun direct impact setup. In the direct impact setup, data from a strain gauge and Digital Image Correlation (DIC) are used to measure the strain and stress wave propagation through the structure. Finally, results from both tests are incorporated into the Machine Learning training data to improve the next prediction.

Gas compressibility in the compression-driven displacement of non-Newtonian fluids in a capillary tube

Kaylah Y. Rhee

Mentors: Sujit Datta, Craig Singiser, and Mohamed Abdelgawad

When a low-viscosity gas displaces a more viscous liquid inside a narrow channel, the interface between them can become unstable and break into finger-like intrusions that reduce how efficiently the liquid is pushed out. Controlling this instability matters for processes ranging from oil recovery to industrial coating. Recent work has shown that the compressibility of the driving gas is a powerful and often overlooked way to stabilize the interface, since a compressible gas naturally softens the injection and delays the onset of fingering. This project investigates how gas compressibility governs displacement in a capillary tube driven by a syringe pump, and extends the idea from simple liquids to non-Newtonian liquids, whose resistance to flow changes with how hard they are pushed.

The displacing behavior is recorded through imaging and compared against theoretical predictions across a range of injection rates and reservoir sizes. By measuring how compressibility interacts with the unusual flow properties of complex liquids, the project aims to clarify the conditions under which such fluids can be displaced smoothly and stably.

Poisson coloring in the hyperbolic plane

Phoebe A. Richter

Mentor: Yassine El Maazouz

Consider a plane with a finite number of points and give each point a certain color. Let points fall in the plane according to a Poisson Point Process and give each point the color of the point in the plane that was closest to it. Aldous showed that in the case of the Euclidean plane, this process results in a random partition of the Euclidean plane into Lebesgue measurable sets. We investigate the process in the hyperbolic plane. We find that the hyperbolic geometry leads to stronger bounds for the distances between particles and their ancestors. We aim to show similar convergence in the plane and investigate any possible stronger properties.

Developing simulation tools for validation of photonic integrated circuits for exoplanet observations through nulling interferometry

Isabelle A. Rivera

Mentors: Nemanja Jovanovic and Anusha J. Pai Asnodkar

Future missions aiming to directly observe Earth-like exoplanets require starlight suppression approaching contrast ratios of 10^{-10} at small inner working angles, exceeding the capabilities of current observatories such as JWST. Nulling interferometry implemented with photonic integrated circuits (PICs) offers a promising approach to achieving these performance requirements while providing compact, stable, and scalable beam-combining architectures. Accurate simulation tools are essential for designing these devices and interpreting experimental measurements.

This work focuses on validating BeamPROP simulations of passive photonic tri couplers by comparison with experimental data collected in the Caltech Exoplanet Technology Laboratory. We develop a Python-based simulation pipeline to automate RSoft simulations while systematically varying geometric and material parameters. Initial studies found that a modeled tricoupler coupling separation of approximately $2.14\ \mu\text{m}$, corresponding to a 7% increase over the nominal design value of $2.00\ \mu\text{m}$, produced substantially better agreement with the measured splitting ratios. Expanding this analysis across all passive tricoupler configurations reduced the overall root mean squared error by approximately 64%, while additional studies preliminarily demonstrate that variations in refractive index contrast may also account for the remaining discrepancies. These results provide insight into potential fabrication errors and guide the design and validation of future astrophotonic integrated circuits for high-contrast exoplanet detection.

Highly scalable atomic force microscope on an unmodified CMOS chip

Katherine C. Robertson

Mentors: Ali Hajimiri, David A. Baum, and Jatin Mathur

Atomic force microscopy (AFM) offers some of the highest-resolution imaging of surface topography among existing microscopy methods. However, existing AFM techniques are frequently limited by throughput, which inhibits their use in, for example, semiconductor manufacturing and quality control. To overcome this limitation, we present massively parallelizable arrays of integrated AFM units. The AFM units are designed to leverage unmodified CMOS processes and maskless post-processing, which enables mass-manufacturability at scale and the co-integration of supporting electronics. Each AFM unit comprises an electrothermally actuated Z-axis stage and cantilevered probe tip. The resonant frequency of the probe tip is read capacitively and used to determine the probe's interaction with a sample surface. As probe-sample interactions occur, the resonant frequency of the beam fluctuates, which provides a measure of the probe's distance to the sample surface. Several non-contact mode AFM units and test structures were taped out on a 180 nm bulk CMOS process.

Investigation of Fis1-mediated regulation of mitochondrial function via proteometabolic profiling

Amarri Z. Robinson

Mentors: David C. Chan and Sangheeta Bhattacharjee

Mitochondria are membrane-bound organelles with vital cellular functions. Beyond the production of chemical energy via oxidative phosphorylation, mitochondria play crucial roles in metabolism, stress response, and apoptosis. Mitochondria are unique in their dynamism: undergoing coordinated cycles of fission, fusion, and self-degradation (mitophagy) to maintain the health of the cell's mitochondrial population. One mitochondrial dynamics protein that has remained elusive is Mitochondrial Fission Protein 1 (Fis1). Bound to the outer mitochondrial membrane, Fis1 recruits numerous protein signalling factors. Overexpression of Fis1 can cause mitochondrial fragmentation in certain cell types, indicating a role in fission. Additionally, Fis1 binds to the TBC1D15 protein, which regulates mitophagy. We aim to investigate the effects of Fis1 gene disruption (knockout) on the mitochondrial proteome and metabolome. This analysis requires effective isolation of the mitochondrial population from a cell lysate. We accomplish this with HA-immunopurification, where HA-tagged mitochondria are isolated using anti-HA antibody conjugated magnetic beads. Following immunopurification, we analyze the proteometabolome of the organelle using Liquid Chromatography- and Ion Chromatography-Mass Spectrometry. This work will provide a higher-resolution view of Fis1's regulatory functions, clarifying its role within mitochondrial dynamics and cellular homeostasis.

Directional Gaussian smoothing for gradient descent

Ryan M. Rodrigue

Mentors: Federico Pasqualotto and R. Michael Alvarez

Traditional gradient descent on non-convex objective functions often gets trapped at local minima or saddle points because local gradients poorly reflect the broader loss landscape. This project develops a gradient descent method that smooths the objective function with a Gaussian distribution, allowing the algorithm to see past shallow local minima and move towards better solutions. This project derives the update equations for isotropic smoothing, using Stein's lemma to obtain a gradient-free estimator for black-box optimization. This method, called Isotropic-Smoothed Gradient Descent (ISGD), was tested on a double-well function with two minima of unequal depth, achieving up to 94% success in selecting the global minimum, compared to 56% for traditional black-box gradient descent. Additionally, initial testing on MNIST data suggests that ISGD shows promise as a complementary method to traditional gradient descent for training neural networks for classification problems. Remaining work extends this framework to anisotropic smoothing, where the smoothing directions and variances adapt to the geometry of the objective, along with a formal convergence proof and experiments on higher-dimensional problems for the anisotropic version.

Limitations of EC-labeled contrastive learning for reaction-enzyme retrieval

Varun E. Rodrigues

Mentors: Frances H. Arnold and Sonia Yuan

Contrastive learning has emerged as a leading computational method for linking enzyme sequences to chemical reactions. However, it has not been assessed whether such a linkage is genuine or how well contrastive models can learn to represent the data. In this work, we compare multimodal models trained on EC-labeled reaction and sequence data to unimodal models trained only on EC-labeled reaction data across a diverse set of reaction encoders. We find that unimodal models convincingly outperform multimodal models, that bespoke reaction encoders have marginal to no benefit over general encoder architectures, and that contrastive models do not greatly improve upon untrained baselines, facing reproducible performance caps on out-of-distribution queries. This leads us to believe that EC-labeled contrastive models do not learn a meaningful linkage between sequence and reaction space. Finally, we hypothesize that EC-labeled contrastive models rely heavily on pattern matching and are fundamentally limited by the curation of the EC system and the model's inability to learn the underlying chemistry.

Tracing radial dust structure in the HD 141569A circumstellar disk

Zihua Rong

Mentors: Konstantin Batygin and Yinuo Han

HD 141569A is surrounded by a prominent gas-rich hybrid debris disk, which brings critical insights into circumstellar dust evolution and planet-forming environments during the transitional phase between protoplanetary and debris disks. To investigate radial dust grain variations across the system, we compare near-infrared scattered light from JWST/NIRCam with mid-infrared thermal emission from JWST/MIRI. Forward-modeling the published NIRCam multi-ring architecture into thermal emission through the MIRI 4QPM coronagraph, accounting for the narrower radial profile in the released code, reveals that the observed outer disk is systematically fainter than predicted. This discrepancy demonstrates that near-infrared scattering grains and mid-infrared thermal emitters follow distinct radial distributions. To explore the disk from an independent perspective free of coronagraphic artifacts, we analyze ground-based VLT/VISIR mid-infrared imaging using joint star-disk MCMC fits. By feeding these empirical ground-based constraints back into JWST/MIRI coronagraphic simulations, we establish a multi facility approach to overcome severe 4QPM residuals and a statistical assertion of the inner-disk emission.

High-quality motion capture for mocap-to-simulation pipeline for powered prosthesis

Brittany M. Rosario

Mentor: Aaron D. Ames

The long term goal of this project is a simulation and testing environment for a powered transfemoral prosthesis, where control policies can be developed and tested before they ever touch hardware. Driving that environment with real human motion requires a motion capture source, and building a capture system we fully control is the main goal of the work. Existing whole-body humanoid pipelines, such as TWIST2, take their motion from a proprietary virtual reality headset, which is close, tied to a fixed tracking volume, and limited in the body points it reports. To replace that input, we are building our own body-worn motion capture based on sparse inertial and ultra-wideband sensing, following recent work that fuses inertial measurement unit with ultra-wideband ranging between sensors to rein in the drift and jitter that limit inertial only tracking. We are assembling custom sensor nodes from bare components, each pairing a six-degree-of-freedom inertial measurement unit and an ultra-wideband radio driven by a microcontroller, bringing them up in stages, and validating live inertial readout and over-the-air communication between the radios. The completed system will reconstruct full-body pose from the combined inertial and ranging data and stream it into a humanoid simulation in place of the virtual reality input, giving the project a portable, low cost motion source and, ultimately, a testing ground for developing prosthesis controllers before they reach hardware.

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.

Long-range, multi-view smartphone demonstrations for robot policy learning

Ryan M. Rudes

Mentors: Soon-Jo Chung and Fengze Xie

Full-body humanoid control research remains bottlenecked by the burden of collecting admissible on-demand motion data. Both traditional motion capture and markerless optical systems are expensive and have a fixed capture volume, making it impossible to capture long-range trajectories. We provide a mobile app bundling iPhone's advanced camera and visual odometry features with calibration tools to allow for sensible and accessible collection of spatially unconstrained, synchronized multi-view demonstrations. We also provide a mechanism for correcting odometric drift using fiducial tags, and offer a more careful execution of fiducial drift correction and acoustic synchronization than seen in prior work. In turn, we demonstrate that conventional devices can provide sufficiently high-quality data for humanoid policy learning.

Zone axis pattern indexing for crystallographic dereplication in autonomous MicroED

Benjamin X. Rudo

Mentors: Hosea Nelson and Dmitry Eremin

Data quality is the cornerstone of successful autonomous data collection and the AI models that increasingly guide it. This holds particularly true in MicroED, where each dataset demands substantial collection time for every microcrystal, so advancements that ensure this time is spent on crystals that merit the investment are especially valuable. Zone axis pattern (ZAP) indexing is one such advancement, enabling rapid crystallographic pre-screening. During initial crystal assessment, ZAP indexing can facilitate dereplication, by comparing diffraction images against previously solved structures or crystals already collected within the same session to avoid unnecessary data collection. For natural product discovery, this same approach can flag whether a crystal resembles a known target, increasing selectivity in which crystals warrant full structural analysis. As part of ongoing development of REyes, our autonomous MicroED platform, we investigated whether SingleCrystal software could serve as a ZAP indexing tool for dereplication, assessing diffraction frames from both still and continuous-rotation datasets. While SingleCrystal successfully indexed a subset of frames from continuous-rotation data, single still images proved substantially less reliable, requiring image-specific parameter adjustment rather than a consistent, automatable protocol. These preliminary results support the feasibility of ZAP indexing for dereplication while highlighting the need for a more robust, customizable approach, motivating development of an in-house ZAP indexing module for direct integration into REyes.

Adapting high-resolution spectroscopy pipelines for IGRINS-2: Studying WASP-189b's atmosphere

Trisha A. Sabadra

Mentors: Fei Dai, Katherine L. Bouman, and Jingyi Zhang

Ground-based high-resolution spectroscopy resolves a planet's spectrum into individual atomic and molecular lines, letting us characterize the composition and dynamics of exoplanet atmospheres. High-resolution cross-correlation spectroscopy has recently begun to resolve the dayside atmospheres of ultra-hot Jupiters in detail, including WASP-189b, whose dayside emission spectrum Sanchez et al. recently characterized with IGRINS, detecting Fe I, Mg I, Si I, H₂O, CO, and OH. A dayside detection alone, however, cannot reveal how heat absorbed on a planet's permanent dayside is transported to and processed on its night side. The principal objective of this project is to determine whether the same species detected on WASP-189b's dayside are also present in its night-side and terminator atmosphere. To do this, we adapt PANCAKE, the high-resolution cross-correlation spectroscopy pipeline introduced by the KPF SURFS-UP Survey for the Keck Planet Finder, to IGRINS-2, the

upgraded near-infrared spectrograph at Gemini North, and obtain night-side emission spectroscopy of WASP-189b (Program GN-2026A-Q-114). We first validate our approach against a known result, using the publicly available pipeline IGRINS_transit to successfully replicate Sanchez et al.'s published dayside Fe I detection of an archival IGRINS eclipse dataset of the same target (GS-2023A-Q-231). Applying our own adapted pipeline to our primary dataset, however, we do not yet recover a convergent signal at WASP-189b's expected orbital velocity for any of eight atomic and molecular species searched, and report a tentative non-detection for the planet's night side.

Temporal acuity and its perceptual consequences in human intracortical microstimulation

Katelyn A. Sadorf

Mentors: Charles Greenspon and Ueli Rutishauser

Sensory feedback for brain-computer interfaces is typically delivered through intracortical microstimulation of somatosensory cortex, with pulse rate as the primary encoding parameter. This project examined whether human perception is also sensitive to the fine-grained timing of individual pulses, independent of average rate, and whether such timing influences the perceived quality of stimulation. Participants with implanted microelectrode arrays performed two psychophysical tasks: a discrimination task comparing temporally regular and irregular pulse trains matched in mean rate, and an intensity-judgment task comparing trains with identical pulse count and duration but different burst grouping. Discrimination of irregular from regular trains was electrode-dependent and, on the most sensitive electrodes, occurred at intervals substantially finer than a published threshold for conscious perception. Independently, burst grouping altered perceived intensity despite no change in total stimulation delivered. These results indicate that temporal structure shapes perception through two dissociable channels, explicit discrimination and perceived intensity, operating independently of pulse rate. Fine timing therefore constitutes an additional, controllable dimension for sensory encoding in neuroprosthetic feedback systems, distinct from and complementary to conventional rate-based approaches.

Engineering carotenoid production in *Yarrowia lipolytica*

Amelia Saffron

Mentors: Gregory Stephanopoulos, Rustem F. Ismagilov, and Michael J. Ream

Abstract withheld from publication at mentor's request.

Improving zero-shot coordination through test-time search over partner models

Jan Pavel P. Safrata

Mentor: Tony Yue Yu

Ad hoc teamwork requires an agent to cooperate with previously unseen partners whose strategies may differ from those encountered during training. I first investigated whether robustness could be learned through Partner-SAM, a training method that exposes an agent to locally adversarial variations of its training partners. Although these perturbations reached their intended behavioral distance and reduced training return as designed, they did not reliably improve coordination with held-out partners across Level-Based Foraging and Overcooked tasks.

I therefore developed a test-time planning method that evaluates each available action by simulating short future trajectories. The agent models its partner from the interaction history, tests candidate actions in the simulator, and overrides its original policy only when the predicted improvement exceeds a safety margin. In a controlled matching game, Bayesian inference over a partner library increased return from the reactive baseline of 2.000 to 7.287, approximately reaching the theoretical identification ceiling of 7.25. On a fully observed Overcooked task, library-based search improved the frozen policy by 28.3 points, with a 95% confidence interval of [26.5, 30.2]. Search using a fixed, non-adaptive partner model also improved performance, while online partner inference provided an additional 6.69-point gain. A conservative override margin substantially improved reliability by preventing changes based on noisy simulations.

The method remained vulnerable to partners outside its model class: confidently incorrect models could make planning harmful. These results indicate that test-time search can substantially improve coordination without retraining the underlying policy, but its success depends on accurately representing the partner. Current work expands the model class by combining learned policies with parameterized behavioral programs.

Design and optimization of adiabatic coupling to focused-ion-beam-fabricated photonic cavities

Kensei Sakamoto

Mentors: Andrei Faraon and Emanuel Green

Rare-earth-ion photonic devices are a promising platform for quantum information processing, but their performance is limited by inefficient optical coupling. In one exemplary platform based on Yb:YVO₄, light is coupled into focused-ion-beam-fabricated photonic cavities via total internal reflection off 45°-angled cuts, achieving efficiencies of only ~20% and motivating the development of improved coupling methods. In this work, we design and optimize an adiabatic coupling scheme between a tapered optical fiber and a tapered waveguide using finitedifference eigenmode (FDE) and finite-difference time-domain (FDTD) simulations. For Yb:YVO₄, the optimized design achieves a simulated coupling efficiency of 99.35%, representing a substantial improvement over the existing approach. We further extend the design to Yb:CaWO₄, a promising alternative host crystal, and to 45°-cut Yb:YVO₄ chips, obtaining similarly high simulated efficiencies of 99.07% and 99.38%, respectively. Finally, we evaluate the sensitivity of the optimized designs to alignment and fabrication errors and identify the tolerances required for maintaining high coupling efficiency. These results establish the adiabatic coupling scheme as a viable high-efficiency interface for rare-earth-ion photonic devices and provide a foundation for future experimental realization.

Radial curriculum learning for implicit neural representation-based reconstruction of non-Cartesian dynamic MRI

Efe Sakarya

Mentors: Shreyas S. Vasanawala, Yisong Yue, Ali Syed, and Terrence Jao

Magnetic resonance imaging (MRI) reconstruction from an undersampled non-Cartesian k-space data remains as a computationally-intense problem, especially for scans that had motion artifacts caused by respiratory and cardiovascular motion or random bodily motion inside the machine requiring high spatial and temporal accuracy while reconstruction. This project develops scan-specific implicit neural representation (INR) models for dynamic contrast-enhanced MRI and introduces a new curriculum-learning-influenced optimization strategy that gradually changes the loss function being utilized in reconstructing the image from low- to high-frequency. By prioritizing lower-resolution structure during early optimization and introducing finer spatial detail later, the proposed method reduces unnecessary computation while improving convergence. Based on the comparison studies with controlled variables such as coil count, sampling factor, and number of epochs, this approach seems to be producing higher-quality final reconstructions while substantially reducing runtime and peak GPU memory usage. Performance is assessed through reconstruction error, spatial sharpness, noise, temporal consistency, motion representation, and computational efficiency across multiple training and regularization configurations. Preliminary results demonstrate that scan-specific curriculum-training-inspired new INR method improves both reconstruction accuracy and resource efficiency in terms of time and memory suggesting that this optimization method can make neural reconstruction more accurate and practical for high-dimensional dynamic MRI.

Searching for successes and failures of the latest high-redshift galaxy models

Maria J. Sanchez Rincon

Mentors: Charles L. Steinhardt and Dimitri P. Mawet

Astronomers' first look at high-redshift galaxies is often in the form of photometry. For decades, astronomers have assumed that the light from high-redshift galaxies is akin to that from local, low-redshift, Milky Way-like galaxies and have used this to infer their properties, leading to inaccurate predictions of stellar mass, redshift, spectral energy distributions, and other

properties. To resolve these inaccuracies, astronomers have developed new models that account for the early-universe environments in which these galaxies formed. This project aims to test how well our latest models predict galactic properties of the most distant objects by pinpointing not only where these new models excel but also where they fail to give us an accurate idea of the object we are observing.

Investigating urinary calcium isotopes in the context of estrogen-related variations in bone health

Abigail W. Sanderson

Mentors: François Tissot, Esther Lahoud-Heilbronner, and Theo J. Tacail

Osteoporosis is a bone disease characterized by decreased bone density, primarily caused by an imbalance in bone remodeling where osteoclasts (bone resorption cells) outpace osteoblasts (bone-forming cells). Calcium isotopes serve as a potential biomarker to assess and monitor bone mineral balance (BMB). Urine isotopic composition ($\delta^{44/42}\text{Ca}_{\text{Urine}}$) provides insight into osteoclast or osteoblast dominance. The calcium isotopic signal varies with diet, exercise, and alcohol consumption. However, the effect of physiological estrogen fluctuations throughout the menstrual cycle has not been investigated. Post-menopausal women are a high-risk population for developing osteoporosis. Following menopause, there is a significant decrease in estrogen levels. Estrogen plays an important role in bone health by increasing the lifespan of osteoblasts. This study seeks to understand how fluctuations in estrogen are associated with changes in urinary calcium isotope composition. The $\delta^{44/42}\text{Ca}_{\text{Urine}}$ values are quantified using a Neptune multicollector inductively coupled plasma mass spectrometer (MC-ICP-MS). Estrone-3-glucuronide (E3G), a urinary metabolite of estrone and marker of estrogen fluctuations, was quantified using the Mira Hormone Monitor. Preliminary results demonstrate successful collection of E3G measurements, and continued monitoring will allow for the evaluation of a potential relationship between estrogen variations and urinary calcium isotope composition.

Modeling vaporization and pre-ignition dynamics in a cubic test cell relevant to the ASTM E659 apparatus

Jorge Sanz Sanchez

Mentors: Joseph E. Shepherd and Charline Fouchier

Autoignition temperatures (AITs) are widely used in hazard evaluations and measured using the ASTM-E659 standard. AIT is not an intrinsic property and depends on the measuring procedure. Prior efforts by the Explosion Dynamics Laboratory at Caltech have examined variability in autoignition testing, but post-injection vaporization and subsequent mixing remain unclear. Spatiotemporal features vary across vaporization regimes, and convective mixing varies across physical configurations, potentially contributing to AIT scatter. This study conducted CFD simulations in a cubic test cell using a reduced-complexity OpenFOAM framework. Initial transient differences among vaporization regimes were suppressed by buoyancy forces arising from the density contrast between air and denser aviation-fuel vapors, leading to similarly stratified fuel-concentration fields across cases. The event also erased the convective characteristics of the pre-injection field. Subsequent fuel mixing was, however, extremely sensitive to wall-thermal conditions, which governed a buoyancy-driven convection process capable of accelerating mixing. This process was also influenced by fuel molecular weight and flask geometry. These findings provide a physical explanation for previously unclear concentration-signal features. The results suggest that differences across vaporization regimes may not dominate AIT variability and highlight the importance of accurately representing wall-temperature boundary conditions in future modeling efforts and testing standards.

Dual-frequency SWV measurements of aptamer sensor drift during extended continuous monitoring in artificial and real sweat

Conor K. Sargent

Mentors: Wei Gao and Hyunah Ahn

Continuous, non-invasive biomarker monitoring could transform cancer detection and management by replacing invasive blood draws with wearable diagnostics. A wearable sensor could let survivors continuously monitor for recurrence, or catch a new cancer at Stage 1 before it progresses. Electrochemical aptamer-based (EAB) sensors suit this well, since aptamers can be engineered to bind nearly any target antigen. This project validates an EAB platform using cortisol as a proof-of-concept biomarker before extending the approach to cancer markers such as PSA and HER2. The gold electrode sensors attached with a cortisol aptamer were tested using a square wave voltammetry frequency sweep (10–60 Hz) to identify the strongest signal pair, then tested continuously (3–24+ hours) in artificial sweat (.25x PBS) and real sweat. Our code corrected or eliminated faulty channels and calculated a dual-frequency ratiometric signal (I_R) to compensate for drift. Over a 24-hour PBS run, I_R remained within approximately 5% of baseline, closely matching published benchmarks. One key limitation was identified, being the fact that the ratiometric signal becomes unstable when either frequency's current crosses zero. With the remaining project time, focus has narrowed to fully validating long-term cortisol stability, with the goal to eventually implement this in cancer detection.

A comparison of computational flow solvers for multi-fidelity aerodynamic shape optimization

Aritz Schube Barriola

Mentors: Steven Brunton, Matthew M. Gherman, Ana Larrañaga, and Patricia Garcia Caspueñas

Designing aerodynamic shapes such as aircraft wings, wind-turbine blades, and vehicle bodies benefits from computational tools that predict how air will flow around a given shape, but these tools face a persistent trade-off: fast tools tend to be approximate, and accurate tools tend to be slow. This project supports the development of a multi-fidelity active learning framework, which aims to combine several such tools of varying cost and accuracy into a single, efficient design-optimization pipeline. The first phase of this work compared four flow-modeling tools spanning a wide range of computational cost and physical detail, from a fast approximate airfoil solver and a neural-network surrogate trained on it, to a real-time simulation framework and a general-purpose flow solver (SU2) capable of resolving complex, three-dimensional turbulent flow. Based on this comparison, SU2 was selected for deeper validation. A repeatable workflow was developed for generating computational meshes around airfoil shapes and configuring SU2 for efficient, accurate simulation, including changes to convergence criteria and solution-acceleration settings that substantially reduced computation time without sacrificing accuracy. This workflow was used to run approximately sixty simulations across two airfoil shapes and four flow conditions, and the results were compared against published wind-tunnel measurements as well as against the faster tools from the first phase. Predictions agreed well with experiment under normal, attached-flow conditions, with accuracy degrading for all tools near stall. In the final phase, two candidate approaches to aerodynamic shape optimization were implemented and compared: a gradient-based method built on SU2's discrete adjoint solver, and an independent, automatically differentiable solver built in JAX. Both approaches successfully reduced the drag of a test airfoil, converging to different final shapes and requiring markedly different computation times, reflecting the differing flow assumptions and optimization strategies of each. These results establish a validated, efficient simulation and optimization workflow and provide concrete guidance on the conditions under which each tool can be trusted, directly supporting the lab's ongoing development of a multi-fidelity active learning framework for aerodynamic design.

Microfluidic approaches to modeling the placenta-brain axis utilizing human organoids

Louise R. Scriven

Mentors: Rowan Karvas, Marianne Bronner, and Christopher Moraes

The placenta plays a central role in fetal development by regulating the exchange of nutrients and oxygen while secreting hormones, growth factors, cytokines, and other signalling molecules that influence the developing brain. Increasing evidence suggests that communication along the placenta-brain axis plays an important role in both normal neurodevelopment and the pathogenesis of neurological disorders; however, the underlying mechanisms remain poorly understood due to the limited access to human development and physiological differences of animal models. This project aimed to develop and establish a 3D printed microfluidic co-culture platform capable of modelling interactions between human trophoblast and cerebral organoids under controlled culture conditions.

A dual-chamber microfluidic device was designed and fabricated using stereolithographic 3D printing, undergoing iterative optimization of both the device design and fabrication process to produce reproducible devices suitable for biological applications. The chip was designed with unidirectional flow driven by gravity, with resistance channels to control flow rate between the compartments. Two compartments are included in this design to accommodate human placental trophoblast organoids and cortical organoids. Methods for chip fabrication, post-processing, and assembly were established, including the development of a sealing strategy to maintain sterile culture conditions while permitting controlled communication between the two organoid compartments. We have pioneered a microfluidic device platform in the Karvas Laboratory with reproducible workflows for the design, fabrication, and use of 3D printed microfluidic co-culture devices, introducing a new experimental capability that can be adapted for future organ-on-a-chip applications.

By providing a more physiologically relevant model of the placenta-brain axis, this platform offers a foundation for future studies investigating how placental signalling influences fetal neurodevelopment. The fabrication workflow established during this project introduces a new experimental capability within the Karvas Laboratory and provides a reproducible approach that can be adapted and expanded in future axes-on-a-chip research. As the platform continues to be developed and validated, it has the potential to support studies of both normal placental function and pregnancy-related disorders, contributing to a better understanding of the role of the placenta in fetal brain development.

Monocular gate detection and localization vision pipeline for autonomous drone racing with adaptive and condition-aware perception

Rhea S. Senan

Mentors: Joel W. Burdick and Jason Siu

Navigating at high speeds, autonomous drone racing calls for localization that produces accurate pose estimations at such speeds for path planning. This project develops a monocular perception pipeline for reliable pose estimation through gate detection with the onboard compute constraints of an A2RL racing drone. The pipeline combines a segmentation model (SAM/knowledge distilled segmenter), a corner detector (convolutional neural network/QuAdGate line segment detector), and a Perspective-n-Point solver with a Kalman filter. This was tested against simulated images on Jetson hardware and a camera stream showing the gate approach. Across many pipeline configurations, gate detection mainly caused inaccuracies, with the best combination reaching below 40 pixels mean corner error. While testing on the camera stream, the four-corner PnP solver exposed specific failures such as an overfitted model, tracker drift, and a three corner pose ambiguity. This solver has a disambiguator that filters out candidates drastically out of range from its prior predicted position. Building on these findings and the recent work of using reinforcement learning to condition parameters using visual features over fixed parameters, future work may explore an image-based policy for adapting the pipeline's gate detection and pose estimation parameters to the changing racing conditions in real time.

Anytime safe reinforcement learning using safe gradient flows in robotic control

Hampus G. Serneke

Mentors: Aaron D. Ames and Pol Mestres

Reinforcement learning is a powerful paradigm for developing autonomous systems in robotics, where a controller is learned by letting an agent interact with an environment. Here, the agent independently explores different actions, and desired behaviour is reinforced with positive rewards. Recently, commercial interest in reinforcement learning and physical AI has surged. However, while RL is performant, there is limited research on the safety of systems trained with it. As humanoids and quadrotors become deployed in environments with close proximity to humans, having theoretical guarantees of safety becomes paramount. In this work, we investigate recent research for achieving anytime safety, i.e. asserting a theoretical guarantee that the RL-agent remains safe throughout the entire training process with arbitrarily high probability, using a discretization of safe gradient flows. We apply the recent theoretical advancements in simulations in Isaac Lab for tasks such as waypoint navigation for quadrotors with obstacles. We also gradually study the effects of practical relaxations of the theory's assumptions, conduct some theoretical investigations to obtain wider bounds on the learning rate, and try PPO-style estimates for the task gradient. Finally, we implement the theory in hardware on a Crazyflie quadrotor and a G1 humanoid.

Advancing ASTM E659 testing: Thermal modeling of ASTM E659 apparatus

Luis Y. Serrano Laguna

Mentors: Joseph E. Shepherd and Charline Fouchier

Autoignition characterization remains a current challenge for the aerospace industry, as safety characterization of jet fuels lies in the autoignition temperature, which is the lowest temperature at which ignition occurs. Currently, the industry employs the standardized ASTM E659 metric for jet fuel autoignition testing. The Explosion Dynamics Laboratory has developed an automated ASTM E659 apparatus. While the system can heat rapidly, it struggles to reach uniform flask surface temperatures. Furthermore, the current system uses passive cooling. Both of these limitations hinder the testing process for consecutive autoignition tests at different temperatures. A new furnace design that incorporates cooling coils into the system has been designed on SolidWorks. A reduced-order model on MATLAB for the current chamber and new furnace has been developed, factoring in PID-controlled heating and cooling due to the cooling coils. Further work will involve finite element modeling on MATLAB and ANSYS to factor in temperature gradients and the development of a more realistic model that reflects the structure of the current chamber and new chamber design.

Task-driven models for neural speech representations

Yourui Shao

Mentors: Francis R. Willett and Yisong Yue

Clear characterizations of how sequences and language are represented in the brain are crucial for developing targeted decoding approaches for speech brain-computer interfaces. Given neural recordings characterizing low-level phonemic and higher-level syntax representations in the brain, we seek to understand their underlying causal mechanisms: *Why are phonemes stored in overlapping subspaces that move non-uniformly toward output? What conditions induce isolated neural syntax representations (and how and why do they decompose)?* Recurrent neural networks (RNNs) can naturalistically model neural computation, with hidden states linear to analogous neural population activity. We apply mechanistic interpretability to RNNs trained to recall phoneme and word sequences, determining the computational parameters and task settings which induce and thus offer biological mechanistic theories for the geometry observed in neural data. We developed several metrics qualifying the representational geometry by summarizing temporal subspace relationships and examined the effects of training conditions such as sequence length, preparation duration, memory size, regularization, and task diversity. Presently, we find that phoneme sequence production geometry is affected by both preparatory period representations and the necessity for model solutions to handle input or preparatory dynamics. Next, we seek to more precisely qualify such preparatory representations and computations, and extend techniques to syntax representations.

Reflector-solver co-training for large language models

Arjun Sharma

Mentors: Yisong Yue and Atharva Sehgal

Large language model performance is influenced by their learned weights and the context provided to them. Existing paradigms for updating model weights during post-training leverage correct rollouts that must be sampled directly from the model. Other work has shown that providing the appropriate context to the model based on external feedback can enable it to succeed on tasks where it ordinarily struggles. In this work, we present a framework where learning to directly solve a task and learning to synthesize helpful context to support subsequent attempts are complementary objectives. When a model cannot solve a task through direct sampling, it is trained to synthesize reflections, based on information about its previous attempts, that raise the quality of subsequent responses. At training time, this helps synthesize successful rollouts from the solver on tasks that it ordinarily struggles in, thus enabling on-policy reinforcement learning on these tasks with any modern paradigm. At inference time, this learned ability helps the model analyze its prior attempts to generate correct responses. In this work, we analyze the design space induced by co-training a model for these two objectives. Compared to rollout-matched training with existing methods, our framework enables learning on extremely difficult problems during training and achieves higher success rates on held-out tasks. By tuning the two complementary abilities of solving tasks and learning to analyze previous experience, our co-training framework enables expanding the capability frontier of the model throughout its lifetime.

Time estimation in chaotic quantum systems

Nyan A. Sharma

Mentors: John P. Preskill and Shreya Vardhan

During the thermalization of a chaotic quantum many-body system, macroscopic observables will equilibrate to steady-state values. Based on this, we expect that there is increased difficulty to precisely estimating time within the system at late times. We can quantify the system's sensitivity to the parameter of time using a quantity called the Fisher Information. Based on thermalization, intuitively, the Fisher information should decay with time. However, we observe that for a subsystem larger than half the overall system, Quantum Fisher Information, or the maximum Fisher Information over all quantum measurements, remains high. We seek to show a complexity-theoretic constraint on accurately estimating time in chaotic quantum many-body systems. More explicitly, by modeling the late-time state as a Haar-random state, we will show that the Fisher Information for time estimation in a large subsystem is exponentially small in any measurement basis of sub-exponential complexity. We make use of concentration of measure techniques for this argument.

Data-driven modeling of molecular state spaces

Shubham Sharma

Mentors: Karthish Manthiram, Luis Burgos, and Sundar Raam Swaminathan

Abstract withheld from publication at mentor's request.

Calculation and analysis of reduced density matrices for quantum period-finding and integer factorization

Karen I. Shekya

Mentor: Marco Bernardi

Quantum period-finding (QPF) is essential to achieving quantum speed-ups to classically difficult problems. The study of the structure of exact and approximate 1-qubit reduced density matrices (1-RDMs) of quantum circuits has revealed a method of determining an unknown period via root-finding on classical machines. We have worked to determine the properties of and improve this algorithm. Simulating a circuit through 3-RDMs and using those to approximate the 1-RDMs in our root-finding algorithm allows for the discovery of periods of the form $(1, 3, 5, 7) \cdot 2^k$ in polynomial time. Testing several methods of approximation has also revealed an optimal approach: by leveraging the entanglement graph of a circuit, one may select which 3-RDMs to reduce into 1-RDMs via partial

tracing with low additional cost. Further, we have rigorously demonstrated polynomial scaling of our algorithm with respect to the number of qubits and found the distribution of the number of attempts taken by our algorithm to factor large semiprimes.

Length generalization in continuous diffusion language models with spectral representations and neural operators

Jucheng Shen

Mentors: Anima Anandkumar, Jiachen Yao, and Robert Joseph George

Continuous diffusion language models generate text by denoising continuous sequence representations, but generalizing beyond training lengths depends on both the representation and the denoising model. This project studies these two stages on six symbolic tasks. At the representation stage, four COSMOS autoencoder variants test how spectral losses and learned frequency attenuation affect fine-tuning. The attenuation variants adapt much faster, while the slope-loss-only variant catches up with longer training. This suggests that the early gap reflects optimization speed more than representational capacity. At the denoising stage, Transformer and Transformer–Fourier Neural Operator (TFNO) score models are compared in the COSMOS conditional-diffusion pipeline. TFNO does not outperform the Transformer on most tasks, and zero padding has mixed effects. An autoregressive control supports this general result, although prefix-TFNO substantially improves one multiplication task. Current experiments test whether attenuation provides a shorter fine-tuning path and whether spectral objectives conflict with reconstruction gradients. Follow-up diffusion experiments with token-aligned ELF embeddings will test whether the adaptation gap and padding behavior persist outside the compressed COSMOS latent space.

Finding supermassive black hole binaries with machine learning in preparation for space-based gravitational wave detection

Yang Shen

Mentors: David O. Cook and Matthew J. Graham

Supermassive black hole binaries (SMBHBs) are promising sources for multi-messenger detection in both gravitational waves (GW) and electromagnetic (EM) emissions. The upcoming Laser Interferometer Space Antenna (LISA) will provide unprecedented opportunities to detect SMBHBs via GWs. However, identifying electromagnetic counterparts of these systems to localize GW observations remains a significant challenge, as the quasar’s stochastic variability often contaminates SMBHB-induced variability. To address this problem, we develop a machine learning (ML) framework to identify SMBHB candidates from the quasar population. Because of limited positive candidates for SMBHBs, we generate simulated SMBHB light curves for the ML model by combining chirping signals produced by orbital motion, quasar variability modeled by a damped random walk, observational noise, and realistic Zwicky Transient Facility (ZTF) observing cadences. We then train an Extreme Gradient Boosting (XGBoost) classifier with handcrafted variability features from simulated light curves as input. Our results show that our classifier effectively distinguishes quasars with chirping signatures from intrinsic quasar variability. This work also highlights the potential of ML as a first-stage selection tool for large time-domain surveys.

SU(2)-equivariant neural networks for magnetic molecules

Erh-Wei Sheng

Mentors: Anima Anandkumar, Beom Seok Kang, Danish Khan, and Yuchao Lin

Equivariant graph neural networks for electronic structure are built from integer-rank tensors, the representations of $SO(3)$. In magnetic molecules, spin-orbit coupling (SOC) ties integer orbital angular momentum to spin $1/2$, so the relevant states carry half-integer angular momentum and transform in representations of $SU(2)$, the double cover of $SO(3)$ — spinors, paired by time reversal into Kramers doublets.

Existing methods regress the Hamiltonian in a fixed orbital basis, append spin as an auxiliary scalar, or impose time reversal on otherwise spinless features. All contract the half-integer degrees of freedom to integer rank, so none propagates a spinor through the network. The obstruction is gauge freedom: a spinor is fixed only up to a phase, and within a Kramers doublet only up to a $U(2)$ rotation, with no smooth global choice.

Rather than fix a gauge, we classify the operations compatible with both $SU(2)$ covariance and Kramers gauge freedom, and obtain message passing in which spinors are transported between sites along a connection built from local eigenbases, with every readout gauge invariant by construction. On a NequIP backbone we compare full spinor propagation (FullSU2) against an ablation that recontracts to invariants at each layer (LocalSU2). For mononuclear Co(II) zero-field-splitting tensors and Dy(III) crystal-field parameters, FullSU2 and LocalSU2 both outperform the $SO(3)$ baseline. Extension to polynuclear systems is limited by data, not formalism.

Computational modeling of nitrogen reduction intermediates on LaFeO_3 and LuFeO_3 perovskite oxide surfaces

Daniella E. Sherman

Mentors: William A. Goddard III and Sejun Kim

This study used density functional theory calculations to investigate nitrogen adsorption and nitrogen reduction behavior on LaFeO_3 and LuFeO_3 perovskite oxide surfaces. The goal of this project was to determine how changing the A-site cation of these perovskite oxides from lanthanum to lutetium affected surface adsorption, nitrogen activation, reaction intermediates, and competition with the hydrogen evolution reaction. First, surface models were constructed and optimized. Nitrogen adsorption was then investigated at multiple surface sites. Nitrogen adsorption calculations were followed by calculations of hydrogenated nitrogen intermediates and adsorbed hydrogen structures. Binding energies were calculated relative to the clean surfaces and reference molecules to compare the varying stability of the different evaluated adsorption geometries. The results of this project showed that nitrogen adsorption depended strongly on both the adsorption site and surface composition. Particularly, LaFeO_3 and LuFeO_3 favored different adsorption geometries and showed differences in adsorption strength and intermediate stability across various calculations. These findings indicate that changing the A-site cation affects the surface environment of perovskite oxides and can provide a future strategy for tuning perovskite oxide catalysts for nitrogen reduction and sustainable ammonia production.

Super-resolution for high pressure quantum sensing with NVs in round-table diamond anvil cells

Ahaan Shetty

Mentors: Norman Yao, Thomas F. Rosenbaum, and Srinivas Mandyam

This project develops a super-resolution optical platform for imaging superconductors under extreme pressure. Diamond anvil cells (DACs) generate high pressures by compressing a sample between opposing diamonds. Nitrogen-vacancy centers embedded near the diamond surface serve as quantum sensors of magnetic fields, temperature, and strain. However, conventional diamond anvils introduce strong optical aberrations through refraction at the air–diamond interface, limiting spatial resolution to several micrometers despite a diffraction limit of a few hundred nanometers. To address this limitation, we employ Diamond Anvils with a Round Table (DART), whose spherical upper surface reduces refraction and aberration by allowing light to enter normal to the interface. Confocal microscopy and initial benchmarking measurements demonstrate improved imaging performance, with resolution approaching the diffraction limit. A wide-field microscope is used to implement structured illumination microscopy (SIM) to improve resolution by approximately a factor of two. We image magnetic vortices in the cuprate superconductor Yttrium barium copper oxide (YBCO), which we use to characterize London penetration depth and relationship with superfluid stiffness. This provides an exciting avenue to probe new physics, as the first pressure-tuned study of superconductors with such high spatial resolution.

Experimental characterization and simulation of nanopore ion beams for single-molecule mass spectrometry

Kensuke Shimojo

Mentor: Michael L. Roukes

This project investigates a nanopore ion field emission platform for high-sensitivity mass spectrometry and its eventual integration with a micro-electrostatic ion trap for single-molecule analysis. While conventional electrospray ionization relies on droplet evaporation, Coulomb fission in a background gas, and a differential vacuum setup, a nanopore source emits ions directly from solution into high vacuum. This approach can drastically reduce ion losses, yielding potential for single molecule mass spectrometry, a powerful tool in proteomics as low-abundance proteins can play essential roles in biology. The short-term goal of this work is to characterize the spatial profile of the unfocused ion beam emitted from a glass capillary nanopore using a microchannel plate detector and to establish a COMSOL model for predicting ion trajectories through electrostatic optics. Progress to date includes setting up a high-vacuum chamber, developing a compact emitter-chamber and load-lock design, and creating a simple COMSOL simulation of a divergent ion beam passing through an Einzel lens. The current simulation explores how lens voltage influences beam focusing and provides a foundation for later models incorporating more realistic ion masses, charge states, kinetic energies, and electrode geometries. These experimental and computational efforts will shape the design of an ion source and trap interface capable of efficiently coupling the nanopore emitter to a single-ion mass sensor. Forthcoming results will be discussed.

A principled approach to mouse motor behavior analysis in Parkinson's disease models

Alyssa H. Shin

Mentors: Sarkis K. Mazmanian, Anastasiya Oguienko, and Justin S. Bois

Parkinson's disease (PD) is a progressive neurodegenerative disorder that impairs motor function via alpha-synuclein aggregation followed by death of dopaminergic neurons. To study the underlying mechanisms of the disease, researchers often use mouse models imitating different aspects of PD. Motor behavior assays are used to assess the severity of symptoms, but behavior is often scored using subjective or simplified rating scales, resulting in one or two data points per trial and limiting reproducibility across studies. This project develops a quantitative tool for characterizing behavior in the hindlimb clasping mouse motor assay using an alpha-synuclein overexpression PD mouse model. Mice were recorded throughout the experiment and pose estimation software was used to track individual body parts over time. From these positional data, I developed measurable characteristics hypothesized to distinguish affected mice from healthy controls. Then I used Bayesian inference to develop hierarchical generative models for these behavioral traits, which allowed the models to account for variability across individual mice and experimental cohorts while estimating differences at the genotype level. This approach offers a more rigorous way of quantifying motor behavior in mouse models and could help standardize how motor deficits are measured across labs studying preclinical models of PD.

First-principles approach to optical excitations of coherent phonons

Han Na Shin

Mentors: Marco Bernardi and Sergei Kliavinek

Recent experimental advances in nonlinear phononics have demonstrated the ability to induce and manipulate properties in materials via coherent phonon excitations. These results highlight the capabilities of engineering practical non-equilibrium functionalities for memory storage and integrated microelectronics, while also playing important roles in chemical and biological activity in molecular systems. However, a comprehensive framework capable of describing, simulating, and predicting the full electron-phonon system in this phenomenon is lacking. We aim to fill in this theoretical gap by developing a first-principles framework that combines the Boltzmann transport equation (BTE), which captures the thermal phonon population, and an equation of motion that encodes the presence of coherence via displacement amplitudes. This framework accounts for phonon-phonon and

electron-phonon interactions, while explicitly retaining both coherent phonon and thermal phonon contributions. Further, we extend this approach to encompass spatiotemporal order arising from parametric resonance, providing a unified first-principles approach to coherent dynamics and their subsequent energy dissipation into thermal phonon and electron populations.

Certifying and learning Hamiltonians with single-qubit gates

Myeongjin Shin

Mentors: John P. Preskill and Yu Tong

Learning an unknown many-body Hamiltonian from real time evolution data is a fundamental problem in quantum science, with applications ranging from benchmarking quantum simulation to characterizing quantum devices. In this work, we study this problem in a highly restricted but experimentally motivated setting, where one can only prepare product states, apply single-qubit gates and measurements, with forward real-time evolution under the unknown Hamiltonian.

We present an efficient algorithm for learning sparse local Hamiltonians without using ancillary qubits, entangling operations, controlled evolution, inverse evolution, or short-time control whose precision depends on the target accuracy. Our algorithm learns the unknown Hamiltonian by using the total evolution time in the Heisenberg limit and also matches all the previous Hamiltonian learning complexities in almost every aspect.

The main idea is to engineer single-qubit gates that transform time evolution under the unknown Hamiltonian into measurable phase differences associated with local Hamiltonian parameters. These phase differences can be estimated with Heisenberg-limited scaling, which approximates the linear measurement data for Hamiltonian reconstruction. Combining this approach with sparse reconstruction and extrapolation techniques yields an efficient learning protocol. Unlike approaches that require increasingly short evolution segments for higher accuracy, our protocol maintains precision and sparsity independent time resolution.

These results establish that state of the art Hamiltonian learning can be achieved using only single-qubit discrete control and measurements, despite the absence of the quantum resources typically assumed in previous approaches. Our work provides a route toward experimentally realistic Hamiltonian learning protocols for quantum devices with limited control capabilities.

Elucidating the molecular mechanisms of DDX19 nucleocytoplasmic trafficking

Katie Shou

Mentors: André Hoelz and Jahaun Azadmanesh

The export of mRNA from the nucleus to the cytoplasm is a vital step in eukaryotic protein production. DDX19, a protein involved in mRNA export, ensures directionality of this process by removing proteins bound to mRNAs as they exit the nucleus at the cytoplasmic face of the nuclear pore. Interestingly, studies have found that DDX19 translocates into the nucleus in response to DNA damage and clears harmful RNA-DNA hybrids. However, the mechanisms through which DDX19 may be transported into the nucleus remain unclear. Additionally, the interplay between its role in mRNA export and its nuclear role, including how DDX19 activity modulators impact its transport and vice versa, is unknown. In this work, we confirm and structurally characterize the interaction between DDX19 and the nuclear transport factor TNPO3 via size-exclusion chromatography with multi-angle light scattering (SEC-MALS) and single-particle cryo-electron microscopy (cryo-EM). We also explore the effect of DDX19 activity modulators on DDX19•TNPO3 complex formation. Our findings reveal that TNPO3 binds to the C-terminal domain of DDX19 through multiple interfaces, including some necessary for activity modulator binding, in a predominantly electrostatics-based interaction. Altogether, these results support that DDX19 nuclear import is mediated by TNPO3 and likely interrupts its cytoplasmic activity, contributing an additional perspective to regulation of mRNA export by the nuclear pore complex.

Fixing nitrogen, faster: A reproducible, high-throughput approach to lithium-mediated ammonia synthesis

Ali Raza Sial

Mentors: Karthish Manthiram and Gangsan Lee

Ammonia synthesis via the Haber-Bosch process consumes about 1–2% of global energy and contributes ~1% of global CO₂ emissions, motivating ambient, electricity-driven alternatives. Lithium-mediated nitrogen reduction (Li-NRR) offers such a route, coupling nitrogen reduction to renewable electricity in a gas diffusion electrode (GDE) system. This project advanced two interdependent objectives: independently reproducing the Lab's Li-NRR results to validate the platform ahead of journal submission, and engineering a high-throughput reaction system to make Bayesian optimization of the reaction's multidimensional parameter space tractable. GDE cells were run under potentiostatic control, with ammonia quantified by ion chromatography and independently cross-validated by UV-Vis spectroscopy. Reproduced Faradaic efficiencies and ammonia production rates matched the original dataset in both magnitude and trend across varying pressure conditions, and agreed between both quantification methods, establishing a trusted experimental baseline and clearing the path to publication. In parallel, we designed and validated a leak-tight, quick-connect cell that maintains stable pressure under prolonged operation, laying the groundwork to scale throughput from 1-2 toward a target of 8 experiments/day. Together, this validated baseline and throughput infrastructure position the project to begin experiments that will feed a Bayesian optimization model, with the goal of identifying optimal operating conditions for ambient Li-NRR.

Reduction pathway control and characterization of HIAM lattices

Ziad Siala

Mentors: Julia R. Greer and Anna Wu

Hydrogel infusion additive manufacturing (HIAM) is a highly modular metal additive manufacturing system that enables high-precision production of metal or ceramic structures. 3D printed hydrogels are infused with metal ions then undergo calcination to produce a metal oxide, which is subsequently reduced to yield a metal structure. This process decouples geometry from material selection, enabling a large range of structure-material combinations. Previous investigations have reported unreduced oxides, porosity, cracking, and significant shrinkage resulting from HIAM thermal treatment. In the first part of this study, a parametric investigation of HIAM reduction was performed to understand how different thermal profiles influence porosity, shrinkage, and cracking. Controlled reduction profiles also enabled synthesis of iron oxide phases; for example, Fe₃O₄ (magnetite), which has applications in electromagnetic interference (EMI) shielding, was obtained using a peak temperature significantly lower than the standard reduction profile. Scanning electron microscopy (SEM), Raman spectroscopy, and energy-dispersive X-ray spectroscopy (EDS) characterized the resulting morphology, oxide phases, and chemical composition. The second part of this study expands previous investigations of HIAM Cu_xNi_{1-x} alloys through focused ion beam (FIB) micropillar compression testing to obtain additional mechanical characterization data. These results will contribute to machine learning models capable of predicting mechanical performance from material-structure combinations.

Sequential quantum circuits mapping between symmetry-protected gapless models

Shreyas Singh

Mentors: Xie Chen and Vibhu Ravindran

Matter at low temperature often exhibits interesting quantum effects by virtue of occupying a low-energy state of its Hamiltonian, which dictates energy interactions between adjacent particles. Consequently, a phase transition between two systems is detectable by continuously interpolating their Hamiltonians and searching for a discontinuous change in some parameter. An alternative (and largely equivalent) method for classifying phases is inspired from quantum computing, wherein two systems are in separate phases if they cannot be connected by a finite-depth circuit of local unitary operators, instead requiring a linear-depth "sequential" circuit. This latter scheme has successfully classified phases among one-dimensional "gapped" models, whose second-lowest energy state remains separated from the lowest as the system size grows. We investigate the potential of the sequential circuit formalism for classifying gapless models, which has thus far remained unclear.

In particular, we construct an invertible sequential circuit mapping between the two gapless models $X+ZZ$ and $ZXZ+ZZ$, as well as a noninvertible circuit mapping between the critical Ising models $X+ZZ$ and $X+YY$, both protected by $\mathbb{Z}_2$ and time reversal symmetries. The existence of sequential circuits performing these mappings, but not of finite-depth circuits, supports the utility of sequential circuits for classifying gapless phases.

Passive interferometric radar ranging

Anvika Singhal

Mentors: Joshua R. Smith, Azita Emami, and Shanti Garman

Abstract withheld from publication at mentor's request.

Modeling and design of scalable qubits through a novel semiclassical framework for nondispersive fluxonium readout

Ram Sivaraman

Mentors: Gil Refael, Jacob Steiner, Nina del Ser, and Yinan Chen

Fast, high-fidelity readout of superconducting qubits is a foundational requirement for scalable quantum computation. In this work, we consider a fluxonium qubit coupled to a readout cavity. Fluxonium is designed to have unevenly spaced energy levels to prevent leakage into non-computational states. Conventional fluxonium readout techniques operate in the dispersive regime, which limits their readout speed. To overcome this, our research explores a nondispersive readout regime, where measurement times are significantly reduced. The challenge is that nondispersive readout creates rapid energy transfer between the qubit and cavity, risking computational leakage and the destruction of quantum non-demolition properties. We aim to find physical design parameters that minimize this leakage. To do so, we investigate semiclassical approximations for the quantum dynamics of the qubit-cavity system. Using non-perturbative two-dimensional Floquet theory, we provide a novel framework for evaluating nondispersive fluxonium dynamics under arbitrarily strong charge and flux drives. This approach circumvents divergences in time-dependent perturbation theory, allowing for the calculation of exact leakage probabilities into non-computational states. Minimizing this leakage provides a foundation for scalable fluxonium qubit design, which will incorporate modeling environmental damping, optimizing drive parameters, comparing the current semiclassical approach with purely quantum numerics, and preserving qubit coherence during rapid measurement.

Structural-functional integration for enhanced detection of subtle epileptogenic lesions: A framework for a multimodal EEG-MRI clinically-informed approach

Saanvi Skanda Subramanian

Mentors: Shreyas S. Vasanawala, Pietro Perona, Bruno Passebon Soares, and Mateus Aragao Esmeraldo

Purpose: Focal cortical dysplasia (FCD) is a leading cause of treatment-resistant epilepsy and one of the most diagnostically challenging diseases in neuroimaging. Even the best deep learning models available to detect it rely on imaging data alone and achieve only suboptimal accuracy. A key insight motivates the approach taken here: experienced neuroradiologists and epilepsy specialists do not search the entire brain blindly. Instead, they integrate electroencephalography (EEG) and clinical findings that pinpoint the region of abnormal activity where seizures originate, and use them to focus and prioritize the lesion search. This study tests whether an automated pipeline can mirror this search pattern by dynamically adjusting the detection threshold of a 3D nnU-Net according to the lesion location based on different localization scales (hemisphere, lobe, and sub-region). This design measures the upper-bound performance improvement that EEG or clinical priors could offer. **Methods:** A 3D nnU-Net was trained on 3D T2-FLAIR images from 175 FCD cases. Anatomical search zones at varying scales were built with SynthSeg, an atlas-based parcellation that runs in seconds per case, far below the roughly eight hours of model training per fold. Rather than applying the fixed 0.5 probability threshold everywhere, the pipeline lowered it inside the localization zone to surface faint sub-threshold peaks and raised it to 0.90 outside the zone to suppress false positives, scoring a case as recovered when any of the top five candidate peaks overlapped the lesion. **Results:** The model detected 86.3% of lesions (mean Dice 0.57), missing 24 cases (13.7%). Restricting a multi-peak search to the lesion

lobe recovered 14 of the 24 misses (58%), which a chance baseline confirms is well above the 1.4 expected from a spatially uninformative search ($P < 0.001$). Most of the benefit came from restriction to a moderate scale, the correct hemisphere or lobe; finer subregional restriction added no significant improvement over the lobe. Conclusion: This study shows that sub-threshold, spatially coherent signal often persists in cases where a standard nnU-Net reports the case as negative. Future work should focus on integrating real EEG and clinical data to improve deep learning detection pipelines.

Interpreting AGN and star-formation diagnostics in 5MUSES: A sampling-aware comparison of PAH equivalent width and optical/infrared variability

Kristiano Sleiman

Mentors: George Helou and Thomas Lai

Infrared-emitting galaxies can be powered by star formation, accretion onto supermassive black holes, or both. Using the 5 Millijansky Unbiased Spitzer Extragalactic Survey, I examined how spectroscopic and time-domain measurements distinguish these processes. Because the 6.2- μm polycyclic aromatic hydrocarbon (PAH) complex traces star formation whereas the 6.2- μm continuum is associated with emission by dust heated by BH accretion, its equivalent width formed the spectroscopic basis of the analysis. I remeasured it by applying a modern twist to long-established local-continuum approaches, combining spectral decomposition with consistent treatment of weak and undetected features. I then placed these measurements within a broader analysis of galaxy properties and multiwavelength variability, drawing on redshift, mid-infrared continuum slope between warmer and cooler dust emission, optical variability from the Zwicky Transient Facility, and infrared variability from the Wide-field Infrared Survey Explorer and its NEOWISE extension. Infrared-variable sources were concentrated among galaxies with weaker PAH emission, consistent with a stronger active-galaxy contribution, although not every galaxy with weak PAH emission varied. Considering these relationships together, while accounting for differences in observational sampling, provides a framework for interpreting what combinations of spectral and variability behavior reveal about the connections and co-evolution of star formation and black-hole accretion.

Optimal pebbling of wedge sums of cycles and kayak paddle graphs

Gavin H. Slocum

Mentor: Matthew M. Gherman

In graph pebbling, pebbles are placed on the vertices of a graph. A pebbling move takes two pebbles from one vertex and places one on a neighboring vertex. The optimal pebbling number is the smallest number of pebbles in a distribution from which every vertex is reachable. We study how optimal pebbling behaves under graph composition, specifically for wedge sums of graphs constructed from cycles and paths. We establish sharp general upper and lower bounds for wedge sums and prove an edge-removal lemma reducing cycle-based pebbling problems to tree-based problems. Using this reduction and results for spider graphs, we determine the optimal pebbling numbers of octopus graphs and wedge sums of several cycles. We then study broken kayak paddle graphs, a family of trees consisting of two branching vertices joined by a path, and derive exact formulas for their optimal pebbling numbers. Furthermore, by reducing ordinary kayak paddle graphs to broken kayak paddle graphs and optimizing over resulting cycle cuts, we obtain an exact formula for the optimal pebbling number of every kayak paddle graph.

Improving the modelling of cement carbonation and carbon uptake via Multiphysics software

Logan A. Smith-Perkins

Mentor: Melany L. Hunt

The production process of concrete is a substantial contributor to carbon emissions; however, due to a set of carbonation reactions, cement also acts as a carbon sink by absorbing atmospheric CO₂. The current estimate of the mass of CO₂ that can be absorbed by cement is 40-50% of the CO₂ emitted in manufacturing cement. However, the model used in this estimate is an approximation with assumptions that are inconsistent with the real world. Our research develops a refined model incorporating more realistic porosity and atmospheric CO₂ concentrations.

By employing a finite element method (FEM) using COMSOL Multiphysics we simulate the CO₂ diffusion throughout a block cement and incorporate the carbonation reactions that are missing in the approximate model. This improved FEM model provides a more accurate assessment of the amount of carbon absorbed through time and matches experimental data more accurately than the approximate model used in global estimates. Ultimately, this research supports the development of sustainable concrete production and utilization strategies to mitigate climate change.

A close look at the dust clump in the Beta Pictoris debris disk

Colin Smith

Mentors: Konstantin Batygin and Yinuo Han

Beta Pictoris is a main sequence star that has been the subject of extensive studies due to its unique clump of dust and CO gas within its debris disk. The nature of this feature, however, remains a longstanding question in the field of planetary dynamics, raising uncertainty regarding how such a feature could be linked to the system's evolutionary history. Existing studies disagree on the proper motion of the clump, split between whether it is moving as a part of the disk, or sitting stationary. As of 2024, a new set of mid-infrared images has been captured of the system using Very Large Telescope's (VLT) VISIR imager. In this project, we reduce the data using chop-nod subtraction and PSF matching, measure the clump's position through chi-squared profile shift fitting and Markov chain Monte Carlo (MCMC), ultimately constraining its proper motion. With these results, we can address one of the most well-studied—yet unresolved—debris disk asymmetries, shaping our understanding of how planetary systems may evolve over time.

Ruling out an intermediate mass progenitor donor in V404 Cygni

Joe F. Smith

Mentors: Kareem J. El-Badry and Cheyanne Shariat

Typical black holes (BHs) in BH low-mass X-ray binaries (BH LMXBs) are understood to not accrete a significant amount of mass over their lifetimes. This assumption can now be tested through binary evolution accounts due to the recent discovery that the BH LMXB V404 Cygni belongs to a hierarchical triple system, which provides an age estimate of 3-5 Gyr via the tertiary. We employed MESA stellar evolution models for V404 Cygni across grids of initial donor masses and orbital periods and robustly examine influences on the system due to the CARB and RVJ magnetic braking prescriptions, non-conservative mass transfer thresholds via β , and Eddington limited versus uncapped retention. It is found that simulations with CARB generally allow higher mass progenitors than RVJ. Of the satisfactory runs, substantial BH growth occurs for moderately non-conservative mass transfer but models with $\beta \geq 0.8$ fail to reproduce V404 Cygni. However, while the simulations demonstrate that higher mass (2.5-3.0 $M_{\odot}$) progenitors can reproduce V404 Cygni's present state, they evolve to these states much earlier than is allowed by the age constraint provided by the tertiary. Therefore, our calculations favor low mass progenitors near 1.5-2.0 $M_{\odot}$.

A standardized CAN bus data acquisition system for a liquid rocket engine ground test stand

Emily M. Snyder

Mentors: Mory Gharib and Jack N. Caldwell

Throttleable liquid rocket engines are essential for precision landing, reusable launch vehicles, and deep-space propulsion, but their wide operating range makes them more prone to combustion and hardware failures than fixed-thrust engines. This project develops the sensing and control electronics for a ground test stand built to hot-fire and characterize a small-scale liquid rocket engine, which does not yet implement throttling, ahead of eventually validating a throttleable flight system. The stand's avionics form a daisy-chained network of standardized boards sharing 12 V power and a CAN bus, including a combined pressure-transducer/thermocouple board with on-chip cold-junction compensation and isothermal ground layout, and a redesigned valve-position-sensor board rebuilt as a single-tier PCB to eliminate cold solder joints that had failed under vibration. Bench testing showed thermocouple channels reading within 0.3-1 °C of an ice -point reference, and fuses tripping at their rated 2 A, addressing a prior short-circuit failure. High-rate sensor data is still routed (though through a newly consolidated connectorized PCB rather than loose wiring) into an existing acquisition unit.

found to sample at only ~ 18 Hz, too slow for load-cell thrust transients; a faster replacement was designed but could not be brought online before this hotfire. Remaining work focuses on full electrical integration ahead of the test.

Mechanistic studies of β -oxidation in saturated dietary fatty acids

Wiktor Sobczynski

Mentors: Johannes Morstein and Binyou Wang

Mitochondrial β -oxidation is one of the fundamental metabolic processes in the cell, converting fatty acids into shorter acyl-CoA units. Even though the mechanism has been thoroughly studied, there are still unanswered questions regarding the β -oxidation of saturated fatty acids and which ones seem to undergo the greatest extent of oxidation. Here, we utilized lipidomic analysis to understand which fatty acids need the carnitine shuttle to enter the mitochondria and which ones get oxidized the most. We treated HeLa cells with deuterated saturated fatty acids and observed their incorporation into acylcarnitines (CARs), phosphatidylcholines (PCs), and triacylglycerols (TAGs). After fatty acid treatment, PCs and TAGs constituted around 80% of the total lipids in the cell. By investigating them, we optimized the work needed to determine plausible results. We observed the 10 most abundant lipid species of PCs and TAGs. We saw that, out of all the fatty acids, arachidic acid was oxidized the most. These findings suggest that shorter fatty acids are incorporated into multiple lipid species in their native form or undergo elongation, whereas longer fatty acids undergo β -oxidation to a greater extent.

Automated dark-field birefringence microscope for investigation of crystallization kinetics of chiral species on magnetized substrates

Alan Song

Mentors: S. Furkan Ozturk and Sean J. Chen

The two enantiomers of a chiral molecule are nonsuperimposable mirror images of each other, related by a parity transformation. Most chirally selective interactions in chemistry, biology, and pharmaceutical synthesis rely on inherently chiral materials to generate and transfer chiral information. In contrast, chirality induced spin selectivity (CISS) is a phenomenon where electron spin is coupled to handedness, allowing for the generation and transmission of chiral information through even achiral substrates.

CISS-like behavior has also been observed in enantioselective crystallization on magnetically polarized substrates, with a bias in the chirality of macroscopic crystals. The origin of the bias must lie in primary nucleation processes, as secondary nucleation of chiral molecules only serves as a means to propagate the CISS interaction of chiral molecules with the magnetized surface to macroscopic scale. Microscopy is a convenient method to investigate the rate kinetics for small systems such as primary nucleation. The stochastic nature of crystallization kinetics requires automation of the imaging over many experiments to precisely constrain rate constants of crystallization for the two enantiomers of a target species.

For this investigation, the primary target species is ribose aminooxazoline, a central RNA precursor, on a magnetite (Fe_3O_4) substrate. This system has previously shown enantiomeric excess during crystallization and is of interest in origin-of-life research. By investigating the rates of primary nucleation of this system's enantiomers, we hope to provide understanding for the propagation of CISS in macroscopic systems, as well as the origin of life's homochirality.

A Lean 4 formalization of categorical deep learning

James Squires

Mentors: Anima Anandkumar and Robert Joseph George

The wide range of increasingly sophisticated deep learning architectures has created a need for a common theoretical basis. This can be done via categorical modeling. We present a Lean 4 formalization for the categorical theory of deep learning. This includes abstract notions of forward pass, backprop, as well as several applications. This code can be used to formally verify compositional

properties of deep learning architectures, as well as providing a convenient, algebraic language for formal specification. We include a formalization of parametric bicategories to capture network parameterization, as well as weighted optics which can be instantiated to obtain backprop. We then demonstrate several applications of the framework. These include wrapping TorchLean forward pass and backprop of MLPs to demonstrate the instantiation, as well as providing a novel, high level categorical model of neural operators in Lean. As the development of new AI architectures continues, the formalization of categorical deep learning will provide a tool to facilitate high-level specification and verification of compositional properties.

Microscaffolds for CAR T manufacturing: Resin design for bioactive hydrogels

Alisha Sri-Ram

Mentors: Julia R. Greer and Steven Chen

Chimeric antigen receptor T-cell (CAR T) therapy is highly effective but slow and difficult to manufacture, in part because current processes do not control how T cells interact with their materials. This project seeks to develop three-dimensionally printed microscaffolds that better retain and activate T cells as a step toward a compact, point-of-care CAR T manufacturing device. We design a microwell scaffold with an internal lattice based on triply periodic minimal surface (TPMS) geometries to increase surface area and confine cells. In parallel, we formulate bioactive hydrogel resins combining gelatin methacrylate (GELMA) and poly(ethylene glycol) diacrylate (PEGDA) with a light-activated initiator and photoabsorber, and systematically vary composition, temperature, and light exposure during printing. These experiments identify an approximately 10%/10% GELMA-PEGDA formulation and a narrow thermal and exposure window that yields mechanically robust, finely patterned microscaffolds on two different bioprinters. We also produce test samples for mechanical measurements and then quantify scaffold stiffness and evaluate T-cell viability, retention, and activation within these microscaffolds, with the long-term goal of shortening CAR T manufacturing timelines and improving therapeutic efficacy in a single integrated device.

A wireless, oxygen-generating macroencapsulation device for high-density islet transplantation in type 1 diabetes

Aruesha Srivastava

Mentors: Siddharth Krishnan, Wei Gao, and Kecheng Wang

Pancreatic islet transplantation has the potential to restore endogenous insulin secretion in patients with type 1 diabetes, but long-term graft survival is limited by immune rejection and insufficient oxygen delivery. Macroencapsulation devices can protect transplanted islets from immune attack, yet their clinical scalability remains constrained by hypoxia, particularly at high cell densities and increased implant thicknesses. This project investigates a wireless, oxygen-generating macroencapsulation system designed to improve oxygen availability while maintaining a compact and retrievable implant architecture. A resonant inductive power-transfer system was developed to wirelessly drive an electrochemical oxygen generator. The receiving antenna, circuit board, and oxygen-generating element were integrated into a functional implant-side assembly. An incubator-compatible stacked transmitter and passive resonator configuration was developed to replace the NeuroLux cage used in initial testing. The optimized six-turn receiving antenna produced sufficient voltage to operate the oxygenator, and the complete integrated system generated visible oxygen bubbles under wireless power. In parallel, human islets were mechanically downsized to approximately 70 μm to reduce oxygen diffusion distance. Preliminary live/dead fluorescence analysis after 24 hours at 1% oxygen showed a lower measured necrotic area in downsized islets than in native islets. A multilayer BioMed Clear Resin-Ecoflex-PDMS coating was also developed to reduce corrosion in culture medium while preserving oxygen generation during short-term testing. These results show the feasibility of combining wireless power delivery, electrochemical oxygen generation, islet-size reduction, and protective encapsulation within a single platform. The ongoing work focuses on biological replication and functional assessment of insulin secretion.

Designing a gas cell for the first ultrafast all-electron spectrometer

Emily A. Stanton

Mentors: Scott K. Cushing and Bethany De Roulet

In this research project, a gas cell was designed and fabricated for an ultrafast all-electron spectrometer that can measure electrochemical processes in a gaseous state, such as CO₂ reduction reaction (CO₂RR). There is currently no method of measuring the intermediates in electron initiated electrochemical processes, which are integral to both further understanding these processes and optimizing materials used in the reactions. An ultrafast all-electron spectrometer will be able to measure this by giving picosecond temporal resolution to instruments at low <10 eV energies. This gas cell has been designed and fabricated, and will be tested under 20 Torr of pressure. The gas cell will also be tested inside the experimental set up with the copper cathode mounted, the CO₂ running through the inlet and outlet lines, and the laser running through the cell.

A pose-based pipeline for automated tennis shot detection and serve analysis for coaching applications

Theodore K. Stoltz

Mentors: Pietro Perona and Ilona Demler

Coaches currently rely on manual video review to evaluate player technique, a process that does not scale to the hours of practice and match footage a team generates and that offers no structured, quantitative record of a player's motion over time. We are developing the CalTennis Coaching Tool, a pipeline that automatically extracts and classifies tennis shots from raw video to support this kind of coaching feedback at scale. Each frame of an uploaded video is passed through RTMPose, a real-time multi-person pose estimation model, to recover 2D keypoints for every visible player, who are then tracked individually across frames. From each player's wrist, elbow, shoulder, and hip keypoints we compute a torso-normalized wrist-velocity signal, and use peak detection together with a set of motion-based rules to identify and classify shot events as serves, groundstrokes, or volleys, with a separate motion-boundary detector giving each serve a precise start and end timestamp. This pipeline has been used to build a large, structured dataset of labeled serve clips and paired pose sequences spanning professional, intermediate, and novice players, which is now being used to train a machine learning classifier to replace the current hand-tuned detection thresholds, with the further goal of testing whether pose-derived motion features can serve as a quantitative marker of player skill level. The pipeline has also been ported to a native iOS app for live, on-device shot tracking, a step toward giving coaches real-time feedback during practice itself.

Probing far-from-equilibrium strangeness-driven bulk viscosity in binary neutron star inspirals

Melvin Storbacka

Mentors: Elias R. Most and Jiaxi Wu

Constraining the high-density equation of state is one of the main goals of gravitational-wave and X-ray observations of neutron stars. Specifically, such observations could help probe the presence of strangeness in the cores of neutron stars. While strangeness is expected to affect static properties of neutron stars, it could also alter their dynamical properties. In particular, if oscillations during inspiral drive strangeness out of equilibrium, significant bulk viscosity can result, which may damp the oscillations. To probe the impact of these dynamical effects, the microphysical evolution of strangeness and beta equilibration must be resolved. To this end, we have developed the first reaction network for strangeness equilibration in neutron-star matter capable of capturing far-from-equilibrium regimes and dynamical coupling between reactions. With this approach, we find that neutron star matter may be driven far from equilibrium during inspiral, giving rise to substantial strangeness-driven bulk viscosity that could affect the orbital evolution and gravitational-wave signal. These findings motivate more careful modeling of nonequilibrium strangeness dynamics during the inspiral.

Extreme long-duration Be star outbursts discovered with NEOWISE

Ming To Suen

Mentor: Lynne Hillenbrand

Be stars are classically defined as B-type stars that exhibit Balmer line emission, indicative of circumstellar gas. They rotate rapidly, ejecting material from their equatorial regions to form a decretion disk. A subset of Be stars undergoes episodic outbursts associated with disk formation and dissipation, typically with amplitudes of 0.5–1 magnitude and durations of less than 1–2 years. We identify a new class of extreme Be star outbursts using NEOWISE data, characterized by larger amplitudes and significantly longer durations, with brightenings of up to 2.5 magnitudes lasting 5–15 years. Multiwavelength light curves from optical and infrared surveys are used to characterize their variability. Most targets become redder as they brighten in the infrared, consistent with circumstellar disk formation, while optical behavior ranges from correlated brightening to anti-correlated dimming, suggesting inclination-dependent effects. Additional evidence for circumstellar material includes H α emission and infrared excess that increases with wavelength. Spectroscopic data, including Gaia XP and ground-based observations, support early-type classifications and the presence of emission-line features. These results suggest that the observed variability is driven by sustained mass injection and long-term disk growth, representing an extreme and previously underexplored regime of Be star activity.

Sub-cycle quantum control of Hubbard excitons in a Mott insulator

John M. Sullivan

Mentors: Matteo Mitrano and Nai-Chang Yeh

Quantum control of strongly correlated systems is a central goal of quantum materials research because it offers the potential to unlock highly advantageous emergent behaviors. We investigate the use of coherent dressing with periodic non-resonant fields to achieve quantum control of bound holon-doublon pairs, called Hubbard excitons, in the one-dimensional Mott insulator Sr₂CuO₃. We utilize a mid-infrared terahertz (THz) pump field in the 4 THz range and sub-picosecond resonant third-harmonic generation (THG) probing to perform measurements, advancing on prior mid-infrared THz work with the aim of inducing Autler-Townes splitting, Rabi oscillations, and coherent remapping of wavefunction trajectories across the Bloch sphere.

Learning to control: A differentiable optimization approach to robust trajectory generation

Priyanka Supraja Balaji

Mentors: Zachary Manchester and Soon-Jo Chung

Robots operating in the real world must adapt to changes that are difficult to capture in advance, including variations in payload, contact, friction, actuator performance, and the surrounding environment. We propose a data-driven trajectory optimization framework for reliable behavior control on hardware. Initially, the controller leverages motion data collected offline from simulation, demonstrations, or previous robot experiments to generate feasible trajectories in real time. During deployment, the robot continuously gathers online measurements to update its internal representation of the system dynamics. Rather than relying on new unconstrained data that can cause the model to drift, the framework uses physics-based filters to ensure incoming measurements are consistent with the expected local dynamics. Simultaneously, differentiable optimization evaluates the sensitivity of the resulting motion to changes in the learned dynamics, allowing the controller to prioritize new data updates that improve task performance. This combination of physical knowledge and task-driven learning facilitates robust online adaptation while mitigating the risk of destabilizing updates. Overall, this work aims to achieve responsive and robust behavior control that refines itself through experience while remaining bounded by physical constraints. We validate the proposed framework through real-time flight experiments on a Crazyflie 2.1 quadrotor subjected to unmodeled aerodynamic disturbances.

Effects of IFN γ signaling on Treg phenotype, transcriptional state, and suppressive function

Amina Tajammal

Mentors: Andrew Luster, Ellen Rothenberg, and Alexis Yero Diaz

IFN γ is a dominant cytokine produced during influenza infection and signals through the IFN γ receptor (IFNGR1/IFNGR2 complex) to activate STAT1-dependent transcriptional programs. Tregs express IFNGR1 and are therefore directly responsive to IFN γ . Prior studies demonstrate that IFN γ signaling promotes T-bet and CXCR3 expression in Tregs, facilitating adaptation to type 1 inflammatory environments. However, sustained IFN γ exposure may reprogram Tregs away from a GATA3⁺ type 2-regulatory phenotype, potentially impairing their ability to control allergic inflammation. Dissecting Treg-intrinsic IFN γ signaling is thus essential to determine whether IFN γ directly drives Treg reprogramming and functional impairment during viral-induced asthma exacerbation. We propose that RVIs drive Treg reprogramming, thereby reducing their capacity to suppress type 2 inflammation. This project aims to understand how viral-induced cytokines, particularly IFN γ , reshape Treg phenotype and potentially immunosuppressive functions, which is critical for identifying therapeutic strategies to restore Treg function, prevent exacerbations, and improve outcomes for patients with asthma.

Mitochondrial inheritance and fertilization-driven asymmetry in the early mouse embryo

Katherine W. Tam

Mentors: Magdalena D. Zernicka-Goetz and Breanna M. McMahon

With global rates of infertility increasing, the source of the earliest patterning and developmental potential of the mammalian embryo continues to be questioned. Despite longstanding beliefs of development starting with largely equivalent blastomeres, growing evidence suggests that asymmetries emerge as early as fertilization, but how these initial differences are generated and propagated through cleavage and ZGA remains unclear. Here, we investigate whether mitochondrial organization, maternal transcriptional organization, and cytoplasmic dynamics contribute to early symmetry breaking in the mouse embryo. With the use of fixed and live imaging of oocytes, zygotes, and cleavage-stage embryos, we are able to track maternal and paternal mitochondrial populations and the distribution of transcripts in response to fertilization. Preliminary live imaging reveals that mitochondria becomes spatially asymmetric in the zygote, thus differentially partitioning between sister blastomeres which compound with future divisions. We further examine how these differences relate to fertilization-induced cytoplasmic flows and the asymmetric distribution of maternal factors, including YAP1 and transcripts associated with the zygotic genome activation. Together, these findings suggest that fertilization initiates intracellular heterogeneities that may provide an early bias from which developmental asymmetry is progressively amplified.

Formal languages and aperiodic spin chains on the boundaries of hyperbolic tilings

Alyna X. Tang

Mentor: Matilde Marcolli

Hyperbolic tilings are regular tessellations of hyperbolic space whose boundary degrees of freedom live on an aperiodic structure known as a conformal quasicrystal. The boundary geometry can be represented by strings encoding the local arrangement of tiles. Previous work developed a categorical framework connecting context-free languages with spin chain sets. We apply this framework to context-free languages that describe boundary configurations in hyperbolic tilings. Within a fixed tiling, we construct morphisms that encode inflation and regional growth, relating boundary languages associated with different initial regions. We also construct morphisms between boundary languages arising from different hyperbolic tilings. We aim to use these constructions to investigate how transformations of the underlying geometry affect the physical properties of the associated spin chain sets.

Detecting short period transiting brown dwarfs in ZTF

Lev Tartakovsky

Mentors: Matthew J. Graham and Sam Whitebook

Conventional stellar models predict that in fully convective stars such as low mass M-dwarfs (MDs) and brown dwarfs (BDs) magnetic braking—main mechanism of angular momentum loss—is weakened, suggesting that BD-MD binaries with orbital period less than 10 days are extremely rare. However, recent discoveries of BD-MD binaries ZTF J2020+5033 (2 hour orbit) and ZTF J1239+8347 (57 minute orbit) at relatively close distances suggest that such systems may be more common than was initially believed. We present a pipeline to detect short period BD-MD binaries in ZTF Variability Archive (ZVAR) in order to measure the occurrence rate of ultrashort period BDs. First, we perform a cut that selects eclipsing binaries within a certain range of orbital period, color, and absolute magnitude. We use an ELLC light curve modeling tool coupled with Markov Chain Monte Carlo algorithm to perform parameter estimation and produce a list of objects whose light curves are consistent with BD-MD systems. After obtaining a list of potential candidates we plan to conduct follow up spectroscopy and radial velocity observations to increase confidence in our candidates.

Tuning counteractions in molecular qubits to illuminate the effect of spin-phonon coupling and covalency on low temperature T₁

Lindsay N. Taylor

Mentors: Ryan G. Hadt and Jonathan Aalto

In recent years, paramagnetic transition metal complexes have emerged as promising candidates for quantum sensing applications due to their small size and structural tunability. By monitoring electron spin decoherence data of the qubit candidates, high spatial and temporal resolution may be gleaned about the local chemical microenvironment, including the dynamics of electric and magnetic fields, the proximity of nearby ions, and the behavior of paramagnetic nuclei. Such applications, however, rely on a robust mechanistic understanding of how molecular structure and environmental factors contribute to spin decoherence. Our focus in this paper was to further understand the contributions of delocalized phonon modes and local molecular covalency, independent of other electronic factors, on spin relaxation. We honed a previous theoretical model for the decomposition spin-orbit coupling along the progression of phonon modes by allowing for reduced assumptions of anisotropy in orbital overlap and spin-relaxation. This was performed by using magnetically dilute single crystal complexes, controlling for phonon modes through counterion selection, to perform variable-temperature variable-field electron paramagnetic resonance spectroscopy to perfectly quantify the spin relaxation tensor. The technique was then applied by varying covalency while holding electronic and vibronic structure approximately constant to similarly detune the contribution of covalency to spin relaxation.

Banding phenomena in rotating viscoelastic polymer solutions

Isaac L. Thompson

Mentors: Sujit Datta, Craig Singiser, and Mohamed Abdelgawad

Rotating horizontal cylinders are important to industrial processes including cement mixing, food processing, and pharmaceutical coating, and prior work on Newtonian fluids in this geometry has characterized how rotational drag, gravity, viscosity, centrifugal effects, surface tension, and inertial instabilities shape the resulting flow. Far less is known about how viscoelastic, non-Newtonian fluids, common in biological and industrial settings such as blood, toothpaste, and polymer solutions, behave under the same conditions. This project involved designing and building a custom rotating cylinder apparatus, featuring a sealed glass tube system with stable RPM control, and high-speed camera visualization, to directly compare Newtonian and viscoelastic fluid behavior. Using a 93% glycerol solution as a Newtonian baseline and a viscoelastic polymer solution as the non-Newtonian test case, we observe that the polymer solution develops axial bands that are absent in the Newtonian case, suggesting the bands originate from elastic rather than purely inertial or centrifugal effects. Kymograph analysis at 20, 100, and 200 RPM shows that band number and spacing depend on rotation rate, while multi-hour timelapses reveal that bands continue to coarsen and merge over time.

Preliminary trials with 100-micron glass beads suspended in the polymer solution show that banding persists in particle-laden systems as well. Future work will use a dimensionless parameter framework spanning cylinder geometry, fluid rheology, and particle properties to systematically map how tube dimensions, fluid type, and particle characteristics govern band formation.

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.

Question and quiz analysis of open STACK questions: Validating diagnostics with real course data

Ernest W. Ting

Mentors: Franca Hoffmann, Jasmine R. Bryant, and Juma Zevick Otieno

STACK, a computer-algebra-based online assessment system deployed within Moodle, generates detailed per-student, per-question, and per-response-tree data on every quiz attempt. This data is rarely turned into actionable insight because existing statistical methods, such as facility and discrimination indices and item response theory, normally require exporting data and analyzing it outside the classroom platform. This project built tools that surface these statistics automatically inside Moodle. First, a plugin reports per-question and per-branch diagnostics directly within a course, and a second uses Moodle's Analytics API to flag struggling students and problematic questions with models trained on indicators such as syntax-error rate and branch usage. These diagnostics were validated using three years of real quiz-response data from a live course. The analysis compared the platform's facility index and a simplified feedback-effect measure against a full item response theory fit and paired significance testing. Additionally, it tracked how diagnostics for the same question shifted across different student cohorts and following instructor revisions. The results showed that the lightweight indicators closely mirrored their statistically fuller counterparts. Furthermore, evidence from revision impacts and instructor interviews confirmed that flagged questions aligned with issues teachers had independently identified. Ultimately, this demonstrated that in-platform diagnostics can reliably guide instructors toward the specific questions and learning gaps most worth addressing.

How pseudouridylation of tmRNA influences *Vibrio*-squid symbiosis

Laura To

Mentors: Edward G. Ruby and Xiaomeng Hu

The mutualistic association between *Vibrio fischeri* and the Hawaiian bobtail squid provides a valuable model for investigating the molecular mechanisms underlying host-microbe interactions. This project examines how pseudouridylation of transfer-messenger RNA (tmRNA) influences bacterial colonization, luminescence, and host immune responses during symbiosis. We compare colonization efficiency and

bioluminescence among *wild-type* (WT) *V. fischeri* and complemented mutant strains ($\Delta ssrA+ssrA$, $\Delta truB+truB$), as well as a *double mutant* ($\Delta truB$ and $\Delta ssrA$ T346A), to determine the contribution of tmRNA modification to successful host colonization. Quantitative RT-PCR is used to measure changes in bacterial gene expression, while hybridization chain reaction fluorescence in situ hybridization (HCR-FISH) is employed to visualize bacterial localization within squid light organ tissues. In addition, outer membrane vesicles (OMVs) are isolated and purified to evaluate their effects on squid colonization and host responses. We hypothesize that disruption of tmRNA pseudouridylation alters the production of unusual RNA species that are recognized by the host receptor RIG-I, resulting in activation of the TBK1 signaling pathway and modulation of host-microbe symbiosis.

Genetic analysis of complete Arctic charr mitogenomes reveals colonization and differentiation patterns in four Arctic lakes

Madison Elle C. Tongco

Mentors: Arnar Pálsson and Jonathan Kirshner

Arctic charr is widely used to study sympatric speciation, differentiating readily into morphs characterized by feeding and spawning behavior as well as morphology. Previous whole-genome studies have revealed genetic differences across and within four Arctic lakes in Iceland ($n=2$) and Norway ($n=2$). I asked whether mitogenome analysis would be able to separate lakes and morphs and if morph behavior could be inferred from gene-specific haplotypes. Firstly, we discovered 22 short nuclear-mitochondrial DNA segments (NUMTs, total 3439 bp), showing relatively low integration of mitochondrial DNA into the nuclear genome, before assembling complete mitogenomes ($n=285$). Whole mitogenome analysis found both Norwegian lakes to be genetically separable, but Icelandic lakes were more closely related, suggesting that charr colonized Icelandic lakes more recently. Morphs within lakes were generally not separable, except in Lake Sirdalsvatnet (Norway) and, to some degree, the four morphs of Lake Thingvallavatn (Iceland). Pairwise F_{st} among Thingvallavatn morphs showed strong intermixing between planktivorous and piscivorous morphs and weak intermixing between large benthic and piscivorous morphs, indicating sex biased gene flow between the morphs. Our results corroborate proposed models of single-morph colonization of lakes and further morph differentiation due to phenotypic plasticity and advantages in local environmental niches.

Machine learning-driven waveform reconstruction and split-off recovery for the Mu2e electromagnetic calorimeter

Ajla Trumic

Mentors: Bertrand Echenard and David G. Hitlin

The Mu2e experiment aims to observe the neutrinoless conversion of a muon to an electron. Its electromagnetic calorimeter is vital for isolating the 104.967 MeV signal from background events, but high-luminosity environments create spatial and temporal pile-ups that degrade energy resolution which current extraction methods struggle to resolve. In this project, I developed a dual-head Temporal Convolutional Network (TCN) to optimize waveform digitization and split-off cluster recovery. Optimized via a composite loss function, the model jointly predicts energy amplitudes and per-tick peak-probability heatmaps. During the peak extraction step, max-pool Non-Maximum Suppression was applied to isolate discrete signal hits by filtering out adjacent, non-maximum values. Results demonstrate that the TCN successfully deconvolves temporal pile-up, offering a robust, automated foundation for split-off reconstruction.

Circulating microbial and vesicle-associated signatures as potential biomarkers of atherosclerotic plaque states

Emily Tu

Mentors: Rustem F. Ismagilov and Roey Lazarovits

Atherosclerotic plaques vary in their composition and behavior as some remain stable, whereas others become inflamed, calcified, rupture-prone, or otherwise high-risk. The underlying mechanisms that drive different plaque states remain incompletely characterized. One possibility is that microbial and vesicle-associated biomarkers contribute to plaque heterogeneity, particularly if circulating material can be traced to the oral microbiome. Periodontal disease and gingival disruption have been

associated with cardiovascular disease and allow microbial products to enter the bloodstream. Once in circulation, microbial cell-free nucleic acids, extracellular vesicles, and bacterial vesicles could link oral microbial communities and vascular disease. This project aims to investigate whether oral microbial biomarkers are detectable in blood and whether they connect to plaque inflammation, calcification, and necrotic core formation. The current approach focuses on establishing a paired oral and blood sampling framework and optimizing methods for detecting cell-free nucleic acids and vesicle-associated material.

Optimizing *Azotobacter vinelandii* as a hydroponic biofertilizer

Tiffany T. Tu

Mentors: Gözde S. Demirer and Chiara Berruto

With approximately 70 million people being born each year, the pressure on the agricultural industry to meet food demands is escalating. In response, farmers have resorted to excessive chemical fertilizer use. This has resulted in the depletion of non-renewable resources, damage of marine systems via eutrophication, and release of potent greenhouse gas emissions. One sustainable strategy to reduce agrochemical inputs to farms while maintaining yields is to harness the power of beneficial plant-microbe interactions. *Azotobacter vinelandii* is a plant growth-promoting rhizobacteria (PGPR) of significant interest due to its innate nitrogenase enzyme that fixes unusable atmospheric N_2 into accessible NH_3 for plants to uptake. In this project, we study and finetune *A. vinelandii*'s ability to fix nitrogen in hydroponic systems with *Brachypodium distachyon* and *Solanum lycopersicum*. We established reliable nitrogen-deprivation phenotypes in both plant species, identified the best-performing *A. vinelandii* strain and growth medium for maximal NH_4^+ production, and optimized the physical hydroponic system (tube volume, powered aeration, volume) for maximal nitrogen fixation and plant benefit.

Self-driving imaging system for fetal brain magnetic resonance imaging (MRI)

Idil A. Turasi

Mentors: Polina Golland and Babak Hassibi

Abnormal fetal brain development must be detected and accurately characterized during the second or third trimesters of pregnancy for effective prenatal intervention. Because standard ultrasound examinations do not provide sufficient contrast and resolution for identifying smaller abnormalities and accurately measuring malformation size and severity, radiological assessment stands to greatly benefit from high-resolution T2-weighted volumetric MRI of the fetal brain. However, uncontrollable fetal motion during the scan precludes the diagnostic utility of the acquired images, necessitating repeated re-acquisition of motion-corrupted images. This project aims to improve a system that tracks fetal head motion and adjusts the next image slice in real time. I modified the existing segmentation model to identify spin-history artifacts separately from the brain, eyes, and background. I also constrained simulated artifacts to the brain region and tested different training weights. The best model achieved an average brain Dice score of 0.865 compared with 0.872 for the original model without artifact segmentation. Future work will use Gaussian regression of the distance from each voxel to the artifact plane and uniform normal sampling to improve artifact recovery.

Symmetry-preserving reconstruction of Higgs boson pairs and a controlled test of pretrained jet representations

Idil Doga Turkmen

Mentor: Harvey B. Newman

Higgs boson pair production directly probes the Higgs self-coupling and the Standard Model Higgs potential. I study events where both Higgs bosons decay to bottom-quark pairs, producing four bottom quarks ($HH \rightarrow 4b$), where large backgrounds and ambiguous jet pairing make reconstruction difficult. I compare cut-based selections, boosted decision trees, dense neural networks, Lorentz Boost Networks that learn from particle four-momenta, and SPA-Net, an attention network that reconstructs the two Higgs bosons regardless of jet order. In the initial simulated study, the mass-aware boosted decision tree performed best; on a larger sample, SPA-Net achieved strong signal-background separation (AUC 0.969) together with Higgs reconstruction. I extend earlier $HH \rightarrow 4b$ studies by testing

whether a pretrained Particle Transformer (ParT), which learns from particles inside each jet, adds useful information to SPA-Net. Direct feature addition hurts performance, while controlled integration removes this harm but yields no practically meaningful gain, showing that pretrained information must match the target analysis. CMS projects evidence-level HH sensitivity of about 3.5σ with 2 ab^{-1} and 4.2σ with 3 ab^{-1} at the High-Luminosity LHC, approaching the 5σ observation threshold. Extending these methods to boosted Higgs decays could increase sensitivity to heavy resonances and other new processes producing Higgs bosons. <https://github.com/idildogaturkmen/hh4b-surf>

Breaking orbital degeneracies of directly imaged companions using KPIC relative radial velocities

Tyler N. Tuttle

Mentors: Dimitri P. Mawet and Clarissa R. Do Ó

Directly imaged substellar companions often occupy orbits with centuries-long periods, making our 10- to 20-year baseline astrometric observations insufficient for accurately determining the orbital architectures of these systems. To further constrain these companions' orbits, this study uses newly reduced high-resolution radial velocities (RVs) derived from spectroscopic observations of the Keck Planetary Imager and Characterizer (KPIC) for a selection of 10 substellar companions. The companion orbits were fit using the Bayesian orbit fitting package Octofitter.jl with joint modeling based on published relative astrometry, stellar RVs, the new KPIC relative RVs, and absolute astrometry derived from the Hipparcos-Gaia Catalogue of Accelerations (HGCA) and the Calibrated Gaia DR2, DR3, and Hipparcos Catalogue (G23H). The addition of these line-of-sight velocities successfully resolves longstanding degeneracies in orbital plane architectures, eliminating unphysical high-eccentricity orbital families. Furthermore, we propose the first-ever orbital solution for the HIP 79098 system ($a \approx 464\text{ AU}$, $P \approx 5321\text{ yr}$) while drastically tightening the posterior distributions for targets such as HD 1160 B and DH Tau B. Overall, this project updates the orbital parameter posterior distributions over the 10 benchmark targets, offering crucial dynamical constraints to inform companion formation models while contributing open-source tools to the scientific community.

Can we study growth arrest in diverse bacteria using electrochemistry?

Emil K. Uddevik

Mentors: Dianne K. Newman and Kemal Demirer

Most bacterial research focuses on actively growing cells, but this often isn't how they are found in nature. In chronic infections, bacteria form structured communities called biofilms, whose cores often lack the nutrients and electron acceptors needed for growth. *Pseudomonas aeruginosa* overcomes this by producing redox-active metabolites (RAMs) called phenazines, which shuttle electrons through the biofilm to distal acceptors such as oxygen. This allows cells to survive but not grow, which we term the maintenance state. Whether this state is unique to *P. aeruginosa*, or a general feature of bacterial survival is unknown. To test this, we used a high-throughput 96-well electrochemical platform to probe RAM-cycling-dependent anaerobic survival in other clinically relevant bacteria, including *Escherichia coli*, *Burkholderia* spp., and *Streptococcus parasanguinis*. The system applies a constant oxidizing potential and measures current, enabling quantitative, mechanistic study of bacterial growth arrest. We first assessed whether each organism could survive reductive stress by cycling RAMs, then whether cells in this state tolerate antibiotics targeting various cellular processes. Overall, we found that RAM-dependent anaerobic survival is not unique to *P. aeruginosa* and confers tolerance to certain antibiotics, making it a clinically relevant state that many pathogens may occupy during infection.

Spatially resolved star formation histories of fast radio burst host galaxies using Keck Cosmic Web Imager

Mariana V. Vale Taveira

Mentors: Vikram Ravi, Kritti Sharma, and Stella K. Ocker

Fast radio bursts (FRBs) are millisecond-duration extragalactic radio transients whose progenitor systems remain uncertain. Although their host galaxies provide important clues to their origins, measurements integrated across an entire galaxy combine distinct stellar populations and may not represent the environment in which a burst source formed. In this work, we use integral field

spectroscopy from the Keck Cosmic Web Imager to study the local environments of twelve FRBs in nearby host galaxies. We developed a spectral analysis pipeline that uses stellar population synthesis and emission-line diagnostics to map stellar ages, metallicities, and recent star formation across each galaxy. These measurements will be used to reconstruct spatially resolved star formation histories and evaluate candidate delay-time distributions, which describe the interval between the formation of a stellar population and the onset of FRB activity. Convolving each galaxy's star formation history with these distributions produces maps of the expected burst rate under different progenitor scenarios. These predictions, when compared to the observed burst locations, provide a direct test of whether FRBs arise promptly from young massive stars or through evolutionary pathways that operate over longer timescales.

Development of a dual-wavelength MSLA platform for spatially controlled photopolymerization of multi-material nanocomposites

Noah H. Van Beurden

Mentors: Julia R. Greer and Kaleb Funk

Polymer nanocomposites consist of a polymer matrix and surrounding nanoscale fillers that provide enhanced properties to the material. This dispersion of nanostructures can increase the strength to weight ratio of the material by an entire order of magnitude compared to the base polymer alone. Conventional fabrication is a two step process: nanostructure synthesis followed by dispersion of the nanostructures within the polymer precursor. However, this method is susceptible to poor binding between the matrix and reinforcing phases, in addition to agglomeration and poor dispersion of the nanostructures at high loadings, which can worsen the mechanical behavior of the composite. To address these fabrication challenges, we propose a new dual wavelength additive manufacturing process that will allow for the formation of the nanostructures to occur simultaneously with the polymer matrix, alleviating the need for preformed nanostructures. To realize this new process, we construct a dual-wavelength MSLA platform. Using 405nm and 530nm light sources and an LCD screen driven by custom software, we can selectively cure the matrix and unformed nanofiller phases of a dual-wavelength sensitive photoresin separately. We believe this will improve the network connectivity and overall mechanical performance of the resultant material.

Accelerating simulations of black hole accretion using time dilation

Grant D. Van Gerpen

Mentor: Elias R. Most

Black holes grow by accreting matter from their environments, but modeling this process across the full range of physically relevant spatial and temporal scales is computationally demanding. In general relativistic magnetohydrodynamics (GRMHD) simulations, higher characteristic speeds near the black hole restrict the global timestep, limiting how long the system can be evolved at fixed computational cost. To address this limitation, we are testing and extending a time-dilation algorithm in the GRMHD code H-AMR. The method slows fluid evolution near the black hole, enabling larger global timesteps while preserving the system's large-scale, time-averaged behavior. Unlike previous discrete implementations for stationary black holes, our analytic formulation allows the dilation factor to vary continuously throughout the domain and extends the method to moving black holes. We have tested this approach using magnetized Bondi-Hoyle-Lyttleton accretion, in which a black hole moves through a uniform medium, and are conducting convergence studies to identify parameter choices that reproduce the expected behavior. We are also generalizing the method to a black hole following a prescribed orbit in a time-dependent metric, providing a step toward accelerated simulations of accretion in systems such as binary black holes.

Human single-cell aging clocks to identify and characterize candidate anti-aging interventions

Sawyer H. VanZanten

Mentors: Eric Sun and Niles A. Pierce

Recent machine learning models use single-cell RNA sequencing (scRNA-seq) data to predict the chronological age of healthy human cells, tissues, and individuals. When these “aging clocks” are applied to data from an individual who has undergone some special anti-aging intervention, they are expected to predict that the individual is younger than they actually are. Here, I use these aging clocks to characterize and classify known anti-aging interventions. I first compile and preprocess a large corpus of scRNA-seq datasets, each consisting of data from donors who have undergone a pro-aging or anti-aging intervention. By applying an aging clock to each anti-aging dataset, I identify the particular genes that are the main contributors to each intervention’s anti-aging effect. This allows me to group the anti-aging interventions into classes, each potentially characterized by a distinct anti-aging mechanism. I further attempt to map each pro-aging intervention to anti-aging interventions in the corpus that appear to have an approximately inverse effect. Finally, by querying the Library of Integrated Network-Based Cellular Signatures (LINCS), a database cataloging the effect of a wide variety of chemical and genetic interventions on gene expression, I identify new interventions that potentially reverse the effect of pro-aging interventions in the corpus.

Quantum calorimetry: NV-center-based AC thermometry towards megabar pressures

Amith Varambally

Mentors: Norman Yao, David Hsieh, and Srinivas Mandyam

Pressure is a powerful tuning parameter for phase transitions, but above roughly 10 GPa heat capacity becomes impractical to measure. Such pressures are typically reached in a diamond anvil cell, whose sample chamber spans only a few tens of microns. Conventional calorimetry wires a heater and thermometer directly to the sample, and that hardware neither fits nor survives compression. Optical techniques avoid the wiring but yield only the ratio of thermal conductivity to heat capacity, in which a transition that raises both quantities can go undetected. This project addresses that limitation by replacing the thermometer itself. Nitrogen-vacancy (NV) centers in diamond exhibit a spin resonance near 2.87 GHz whose frequency shifts with temperature at approximately -70 kHz/K. Nanodiamonds hosting these centers can be deposited directly onto a sample prior to loading and are expected to remain functional towards megabar pressures. Interferometric pulse sequences adapted from NV magnetometry resolve kilohertz-to-megahertz oscillations of this resonance, the range required for AC calorimetry. We modulate a heating laser and record the resulting oscillation of the NV resonance. Its amplitude determines the temperature oscillation, and therefore the heat capacity. Its phase lag relative to the drive determines the thermal conductivity, so that the two quantities are obtained independently rather than as a ratio. Over the summer we constructed the confocal microscope for this measurement and implemented the pulse sequences that perform the thermometry. These include a variant of the Ramsey sequence in which the magnetic contribution to the accumulated spin phase cancels while the thermal contribution is retained. As initial demonstrations we target the antiferromagnetic-to-paramagnetic transition of chromium near 311 K. We then map the heat capacity across the $\alpha \rightarrow \epsilon$ structural transition of iron from 0 to 20 GPa, alongside magnetization, thermal conductivity, and spin relaxation on the same platform. Together these establish a route to heat-capacity measurement in a regime where it has been inaccessible, and one that extends in principle to megabar pressures.

Actuated perception for sage autonomous target following on humanoid robots

Alexander F. Vasquez

Mentor: Aaron D. Ames

Humanoid robots performing navigation and obstacle avoidance typically rely on perception that is either external and prepared ahead of time, or fixed to the robot’s body, limiting them to a single region of interest at a time. In this work, we equip a Unitree G1 humanoid with a 2-DOF pan-tilt head that decouples camera orientation from body orientation, and evaluate it on two tasks that require continuous, actuated perception of a moving target. In the first task, the robot must track a moving

ball with the head while an MPC and a CBF-QP safety filter govern body navigation around obstacles, using a color-based (HSV) and neural (YOLOv11n, YOLOE-26) detector pipeline evaluated for both detection and onboard inference latency. In the second task, we build on a prior kick-imitation policy and design a policy that decides when and where to kick, improving the direction of the baseline kick. Together with an interception planner, this achieves a continuous wall-rally kick. Using only onboard perception, the same control loop sustains a rally of 8 kicks, limited mainly by motion-induced sensing noise. These results show that decoupling perception from body orientation through an actuated head lets a humanoid keep a target in view and act on it, across both a safety-filtered navigation task and a precision manipulation task, without relying on external infrastructure or a body-fixed viewpoint. Ongoing work includes improving the kicking policy with adversarial motion priors, adding onboard odometry to remove the remaining ground-truth dependency, and validating both tasks on the physical pan-tilt head.

Murasugi sum and the constituent structure of strongly quasipositive knots

Ostap Vyshnevskiy

Mentors: Sergei G. Gukov, Josef Svoboda, and Lara San Martin

Strongly quasipositive (SQP) knots are an important class of knots in low-dimensional topology and contact geometry. Each SQP knot bounds a quasipositive Seifert surface, and understanding how these surfaces decompose is central to understanding the structure of SQP knots. One natural operation on Seifert surfaces is the Murasugi sum, which combines two minimal-genus Seifert surfaces into another. For SQP links, every Murasugi sum can be realized as braided plumbing of their minimal-index braided Seifert surfaces. This project investigates the Grothendieck group of strongly quasipositive links under Murasugi sum, with an emphasis on braided plumbing. A primary motivation comes from the knot invariant F_K , whose leading term is conjectured to behave multiplicatively under Murasugi sum. Identifying a collection of generating SQP links and understanding F_K on these pieces could lead to a systematic description of the invariant for all SQP links. More broadly, this work aims to clarify the algebraic and geometric structure underlying strongly quasipositive knots and links.

Fault-tolerant coordination for the Near-Earth Object Surveyor ground data system

Ahmad Wahab

Mentors: Joseph Masiero and Analetta Ehler

The Near-Earth Object Surveyor mission requires a ground data system that can continue operating despite computer failures, restarts, and network disruptions. Its cluster consists of a leader node that coordinates processing and worker nodes that execute data-processing tasks. Previously, worker membership depended on manual configuration or remote procedure calls, while coordination relied on a fixed leader, creating a single point of failure. This project developed a fault-tolerant coordination layer in Rust that replaces those mechanisms with automatic membership management and leader selection. Membership is maintained using the Scalable Weakly-consistent Infection-style Process Group Membership protocol, in which computers probe one another and spread status updates through randomized gossip. Direct and indirect probes, suspicion timers, and Lifeguard-inspired local health awareness reduce false failure reports during slowdowns or overload. Periodic full-state exchanges repair missed updates, while process-generation and incarnation numbers prevent outdated messages from reviving failed or replaced processes. Leadership is assigned through a PostgreSQL lease, and each new term receives a higher fencing epoch so workers reject delayed commands from former leaders. An elected leader must restore the processing state before accepting work. Tests demonstrate cluster formation, failure detection, restart recovery, view repair, graceful departure, and automatic leadership transfer after shutdown or failure.

A survey on relational database course pedagogy

Robert R. Walker

Mentor: Ethan Ordentlich

Unlike some Computer Science courses that may come in one or two forms, relational database courses have many freedoms that every institute can take in their own version. Caltech's version of such a course is CS 121. Since we are only one of the many schools that teach relational databases, what is the "right" way to implement a relational database course? To do this, we inspected industry leaders in Computer Science such as CMU, Berkeley, UChicago, etc. With a deeper look into these universities and why their courses are designed the way they are, we hope to better understand the pedagogically "correct" way to implement a relational database class and, if necessary, apply these changes to our course to improve our student's academic and professional careers.

Opposing regulation of androgen receptor stability by RNF113A and PSMD14

Anthony H. Wang

Mentors: Jonathan Tsai, Lea A. Goentoro, and Wallace Wong

Uncontrolled androgen receptor (AR) signaling is a major driver of prostate cancer. Although androgen deprivation therapy initially suppresses advanced disease, nearly all patients with metastatic prostate cancer eventually develop castration-resistant disease. Defining mechanisms that regulate AR protein stability may reveal new therapeutic opportunities. Unbiased genetic screens identified the E3 ligase RNF113A as a driver of AR proteasomal degradation, and PSMD14, a deubiquitinating proteasomal subunit, as a stabilizer of agonist-bound AR. In purified-protein pulldown assays, androgen-response element DNA reduced AR-RNF113A binding, suggesting that RNF113A recognizes the AR DNA-binding domain (DBD). AlphaFold 3 modeling predicted specific RNF113A-contacting residues within the AR DBD. Mutating these residues prevented RNF113A binding, whereas transplanting them onto the estrogen receptor increased RNF113A binding, further localizing the recognition site to the AR DBD. To evaluate PSMD14-mediated regulation, we measured AR ubiquitination following chemical inhibition of PSMD14, which restored agonist-bound AR ubiquitination and degradation. We then combined endogenous PSMD14 knockdown with overexpression of PSMD14 point mutants. Flow cytometric analysis of an AR degradation reporter identified candidate residues required for receptor stabilization. Together, these findings support an androgen-dependent RNF113A-PSMD14 axis that regulates AR ubiquitination and stability.

Development of thin-film tactile sensors using PVDF and P(VDF-TrFE-CFE)

Chloe Wang

Mentors: Ioannis Kyriassis, Wei Gao, Eric Chang, and Peter Ballentine

Static and dynamic tactile sensors, which mirror the slow-adapting and fast-adapting mechanoreceptors of the human tactile system, are essential for providing rich sensory information to robotic manipulators. Thin-films are especially well-suited for fabricating tactile sensors as they produce compact and lightweight devices with high sensitivity and accuracy. In this project, we develop thin-film tactile sensors based on polyvinylidene fluoride (PVDF) and poly(vinylidene fluoride-trifluoroethylene-chlorofluoroethylene) (P(VDF-TrFE-CFE)). Chromium-gold electrodes are deposited via electron-beam evaporation through shadow masks to fabricate parallel-plate capacitive sensors on each film. The devices are then characterized under static and dynamic compressive loading by measuring both capacitive and piezoelectric charge response. Charge generated by the sensors is measured using an ADC-based charge readout, while capacitance changes are measured using a CDC. Although PVDF is widely used for dynamic tactile sensing through its piezoelectric response, we investigate if capacitive measurements can also enable static tactile sensing within the same device. We further evaluate P(VDF-TrFE-CFE) for a strong static response through its high dielectric constant. Finally, we aim to separately integrate each device into a compact robotic finger and demonstrate static and dynamic signal acquisition.

Distinguishing gap transients from Type Ia and II supernovae using photometric classification

Harold Wang

Mentors: Mansi M. Kasliwal and K-Ryan Hinds

Gap transients are rare events, dimmer than supernovae and brighter than novae. They consist of luminous red novae, thought to be stellar mergers, intermediate-luminosity red transients (ILRTs), which may be due to electron capture supernovae, and other subclasses. Subclass assignment can be difficult. With the advent of Rubin Observatory's Legacy Survey of Space and Time (LSST), many more gap transients will be discovered, making a better classification increasingly important and requiring reliable differentiation between gap transients and frequently detected supernovae. We collected photometric data – measurements of an object's brightness through filters – for a sample of literature and Zwicky Transient Facility gap transients. Using Gaussian processes, we interpolate the photometric data and resample rare subclasses to generate more lightcurves. For comparison with supernovae, we simulate a sample of type Ia and type II lightcurves. We then use two methods: training a Random Forest classifier on photometric features, and transforming lightcurves into 2D histograms for classification with convolutional neural networks. This work explores machine learning methods that may be useful for LSST classification of gap transients.

A framework for inference in spatial transcriptomics models

Jenna Wang

Mentors: Lior S. Pachter and Kayla C. Jackson

Spatial transcriptomics tools measure gene expression while preserving the spatial organization of tissue. While these tools enable the interrogation of diverse biological systems, current analyses fail to address technical processes that introduce spatial bias and confound biological signals. Here, we have implemented a physically motivated model of spatial gene expression in a modular and interoperable Python module. At its core, the package fits a model that explicitly separates biological signals from technical artifacts while encoding assumptions about underlying gene regulatory mechanisms. The package is validated on simulated datasets with known parameters to confirm accurate recovery of the biological parameters. All features of the package are designed with modular architecture that enables future extensions, including improved estimation of variance components and joint analysis of multiple genes.

Science-aware guidance trajectory design for remote sensing of Antarctic ice streams using a high-altitude pseudo-satellite drone

Jeremy C. Wang

Mentors: Brent Minchew and Evan Kramer

Accurately understanding the dynamics of Antarctica's grounding zone, the region where grounded ice meets floating ice, is critical for constraining uncertainties associated with future sea level rise predictions. The coupled interactions between ice deformation, viscous flow, and water movement at the grounding zone create a highly complex system that current models are unable to capture, because we lack measurements to inform parameter values and model constraints. Specifically, understanding how the grounding zone migrates from tidal forcings requires high frequency measurements that current satellite remote sensing cannot acquire. This project develops a science-aware framework for designing remote-sensing surveys of Antarctic grounding zones using a high-altitude pseudo-satellite (HAPS) equipped with L-band synthetic aperture radar (SAR), focusing on the Thwaites Glacier region in West Antarctica. We identify the key scientific processes and signals that the survey must capture, while accounting for sensing requirements and constraints imposed by HAPS vehicle dynamics. These include the spatial extent and timescales of tidally driven ice deformation, radar measurement uncertainty, and the operational limitations of the HAPS platform. This work provides a foundation for developing optimized HAPS flight trajectories for high science value observations of Antarctic ice streams.

Perceptive terrain traversal for humanoid robots via route-conditioned reference generation and learned tracking

Minran Wang

Mentors: Soon-Jo Chung and Vrushabh Zinage

Humanoid robots require terrain-aware motion plans that are both geometrically feasible and dynamically trackable, especially when traversal involves contact-rich behaviors such as stepping, jumping, climbing, vaulting, or crawling. This work proposes a three-layer framework for perceptive humanoid terrain traversal that separates route feasibility, whole-body reference generation, and physics-based tracking. First, fused LiDAR and camera observations are converted into a labeled terrain graph whose edges encode traversal modes rather than explicit footsteps or velocity commands. Second, a PARC-style reference generator produces dense future whole-body motion and full-body contact labels conditioned on the selected route, terrain map, and motion history. The generator is trained from a physics-corrected motion library and supervised with route, contact, penetration, and forward kinematic consistency losses. Third, a history-aware reinforcement learning tracker follows the generated future reference in simulation while accounting for contact feasibility and tracking error. This formulation avoids hand-specifying footstep timing, enables multiple traversal modes to be represented through route labels, and provides a deployment pathway in which a compact generator proposes diverse candidate references while a learned feasibility model and tracker select dynamically executable motions.

Determining the chaperone activity and interaction interface of *Chlamydomonas reinhardtii* chloroplast signal recognition particle to its client proteins

Muping Wang

Mentors: Shu-ou Shan and Yelim Yi

The 43-kDa subunit of the chloroplast signal recognition particle (cpSRP43) is a chaperone protein that protects proteins relevant to the assembly of the light-harvesting complex (LHC) in chloroplasts. Past research in *Arabidopsis thaliana* (*A. thaliana*), a model plant, has demonstrated that *A. thaliana* cpSRP43 is a dual-function chaperone that both protects light-harvesting chlorophyll *a/b*-binding proteins (LHCPs) in transit to the thylakoid membrane and protects GENOMES UNCOUPLED 4 (GUN4) in the tetrapyrrole biosynthesis (TBS) pathway from heat stress. However, the chaperone activity of cpSRP43 is not well-understood in algae. Here, we demonstrate that cpSRP43 in *Chlamydomonas reinhardtii* (*C. reinhardtii*), a model alga, has GUN4-specific chaperone activity. Via light scattering assays, we show that GUN4 is disaggregated by *C. reinhardtii* cpSRP43 in a concentration dependent manner. We also show *C. reinhardtii* cpSRP43 has weak chaperone activity against LHC5, a representative LHCP and weak binding to the L11 peptide, a conserved cpSRP43 recognition sequence in LHCPs. Our cross-linking coupled with mass spectrometric studies further identified the binding interface of *C. reinhardtii* cpSRP43 to GUN4. This information can be used to engineer *C. reinhardtii* cpSRP43 variants with alternate chaperone activity to modify the assembly of the LHC, and in turn, improve the photosynthetic efficiency of *C. reinhardtii*.

Coarse-to-fine 3D MRI reconstruction via resolution-agnostic neural operators

Ruibo Wang

Mentors: Anima Anandkumar, Jiayun Wang, and Valentin Duruisseaux

Magnetic resonance imaging produces detailed images of the body but is slow to acquire, so a central problem is reconstructing high-quality images from undersampled measurements. We tackle this problem in three dimensions. Reconstructing a whole volume, rather than stacking independently reconstructed slices, keeps anatomy consistent through the slice direction, but it is memory-heavy, and training a network directly at clinical resolution is usually infeasible. Our method uses neural operators and is resolution-agnostic: we train it at low resolution and then reconstruct at higher resolution without retraining, a coarse-to-fine scheme that keeps training memory low. On thirty-three test volumes it outperforms a strong three-dimensional convolutional baseline by about three decibels in peak signal-to-noise ratio. The improvement holds on every volume, persists when measurement

noise is added, and comes with roughly half the training memory. These results show that one resolution-agnostic operator can reconstruct volumetric scans more accurately than a fixed-resolution network while being cheaper to train. We are now extending the evaluation to full-resolution volumes and re-checking every number we report.

Investigating the growth characteristics and extracellular polymeric substances (EPS) of methylotrophic methanogens in AOM consortia

Wendy Wang

Mentors: Victoria J. Orphan and Yurina Takahagi

ANME-SRB live in diverse microbial communities that anaerobically oxidize methane (AOM), syntrophically consuming methane and producing sulfide. Even in long-term laboratory sediment-free incubations, other members of the microbial community, such as methanogens, persist for many years. Existing work in the lab has isolated strains CA12 and MCC4, of the genus *Methanococcoides*, from sediment free cultures dominated by ANME-2 and SEEP SRB-1. In this project we sought to characterize the growth rates strains CA12 and MCC4 using the Microbial Growth-rate Intervalometer (MAGI), develop a protocol for the efficient extraction of EPS while maintaining cell integrity, and investigate the components of the EPS. EPS is a crucial interface between microbes in natural consortia, possibly giving insight into the persistence of less common community members. We quantified nucleic acid, protein and polysaccharide content of the EPS, investigated the spatial distribution of nucleic acids in CA12 aggregates through methyl green-pyronin staining, and identified methods of CA12 disaggregation that may also apply to ANME-SRB aggregates.

Online adaptation for general humanoid whole-body tracking

Yuyang Wang

Mentors: Soon-Jo Chung and Vrushabh Zinage

Humanoid robots must follow complex whole-body motions while their physical conditions change after deployment. This project develops a general motion controller that can recognize and respond to such changes without retraining its neural network. The controller compares predicted motion with sensor measurements after every control step and uses the difference to estimate the robot's current physical condition. This estimate then guides subsequent actions. Experiments show that the estimated condition captures which motor has weakened and how much strength was lost. Initial tests include physical X2 Ultra experiments with changes in knee strength, together with simulated payload and wind-like external forces. These results suggest that online physical-condition estimation can complement a working general motion tracker and improve its ability to operate beyond controlled laboratory settings. More broadly, interactions such as picking up a box or opening a door also change the forces acting on the robot. The framework could therefore, in principle, treat these interactions in the same way as payloads, wind, motor faults, and other real-world disturbances. Future work will test whether this inferred condition improves closed-loop tracking and recovery during diverse physical tasks.

FROST: Stochastic prediction-space adaptation for sim-to-real PDE forecasting

Zichang W. Wang

Mentor: Yisong Yue

Simulation pretraining followed by real-data fine-tuning is a natural strategy for adapting learned PDE forecasters, such as neural operators, to physical experiments. Yet deterministic adaptation can leave a structured mismatch between forecasts and experimental futures, whose variability reflects unresolved dynamics and measurement noise. To this end, we introduce FROST (Frozen-Operator Stochastic Transport), which freezes an upstream forecaster during correction and learns the conditional distribution of its residual from real measurements. Conditioned on observed history and the base forecast, FROST corrects systematic bias while modeling structured variability around the forecast. Across five RealPDEBench systems, FROST improves long-horizon ensemble-mean accuracy and fair CRPS over the corresponding fine-tuned FNO on every system, with reductions of up to 61%

and 75%, respectively, and matches or outperforms every deterministic and generative baseline in mean rollout accuracy on all five systems. Additionally, we find that when applying FROST, the base FNO need not be fine-tuned at all. Together, these results establish stochastic prediction-space correction as a powerful complementary axis of sim-to-real adaptation for PDE forecasting.

Parameter-conditioned neural operators for size-generalizable dynamics in driven Hubbard–Holstein systems

Miles M. Waugh

Mentors: Anima Anandkumar and Chuwei Wang

Simulating the dynamics of strongly correlated quantum materials is computationally challenging due to the exponential growth of the Hilbert space with system size, confining exact simulations to lattices far smaller than those studied experimentally. Fourier neural operators (FNOs), originally developed for solving partial differential equations, naturally accommodate inputs of varying system sizes while capturing nonlocal correlations. Here, we develop a parameter-conditioned Fourier neural operator, trained only on small lattices generated by a hybrid solver that treats electrons exactly and phonons variationally, to predict the time evolution of physical observables in driven Hubbard–Holstein systems at sizes beyond the reach of conventional numerical methods. This approach aims to provide a framework for finite-size extrapolation in strongly correlated quantum materials that generalizes across Hamiltonian and pulse parameters.

Nickel-catalyzed cross-electrophile coupling facilitated by a Co(II) co-catalyst

David V. Welt

Mentors: Nilay Hazari, Theodor Agapie, and Jasmine L. Chow

Nickel-catalyzed cross-electrophile coupling (XEC) is an important method for the formation of C(sp²)-C(sp³) bonds, present in many pharmaceutical targets. To improve control over catalysis, prior work has used co-catalysts to generate the alkyl radical that can then be captured by the nickel species and coupled with the aryl substrate. Since the nickel and co-catalyst cycles are relatively independent of each other, such systems can be optimized by tuning the catalyst loadings to match the rates of intermediate steps. In this project, we aim to use a Co(salen-Me) co-catalyst to couple aryl halides with secondary alkyl chlorides. We have studied the substrate scope, optimized the catalyst system for a variety of aryl and alkyl halide substrates, and quantified yields by ¹H-NMR spectroscopy. Additional work has been done towards obtaining isolated yields and to synthesize hypothesized intermediate species for use in mechanistic studies.

Developing a molecular replacement method for phasing low-resolution MicroED data of paramagnetic transition metal complexes

Kyler E. Williams

Mentors: Hosea Nelson and Samuel J. Mussetter

Existing spectroscopic techniques struggle to characterize catalytically relevant paramagnetic transition metal complexes, and crystallography remains the gold standard. Microcrystal electron diffraction (MicroED) extends this to submicron crystals, but dynamical scattering, poor crystallinity, low completeness, preferred orientation, and short-lived intermediates frequently complicate or preclude ab initio phasing. Here we develop a fragment-based molecular replacement (MR) method for small-molecule transition metal complexes, enabling phasing of otherwise intractable MicroED data. MR searches rotational and translational space for the fragment placement that best reproduces the experimental diffraction data, and its ability to phase beyond the resolution regime accessible to direct methods has made it routine in macromolecular crystallography. To adapt MR to small molecules, we assemble a library of metal-ligand fragment search models, prioritized using the known synthetic route and complementary characterization data, and we have built a Phaser pipeline suitable for transition metal complexes. We demonstrate the approach on iron bipyridine complexes, including structures in which Fe(II) and Fe(III) species co-crystallize within a single asymmetric unit. Ultimately we aim to deliver an automated, user-directed workflow for routine structural characterization, with broad applications in catalysis and high-throughput structure determination.

SLAC ATLAS ITk pixel testing and beam-induced background in the FCC

Margaux H. Wong

Mentors: Charles Young and Harvey B. Newman

Over the next five years, the Large Hadron Collider (LHC) at CERN will be upgraded to increase the luminosity and thus probability of observing rare interactions such as Higgs boson self-coupling. SLAC is assembling the Inner System of the new ATLAS Inner Tracker to be shipped to CERN in final form. During module reception, pixel modules were visually inspected and output data integrity, front-end chip functionality, sensor longevity, and bump bond connectivity were tested. Modules withstanding transit conditions during final shipment undamaged is critical. Thus, a test was designed to study module shock tolerances with results determining shipping procedures. The Drop Test setup was optimized to produce consistent acceleration over multiple iterations simulating measured shipping parameters. Modules were dropped with an increasing drop count and subsequent performance tested. Separately, theoretical work was performed on the Future Circular Collider, a possible successor to the LHC designed to enable new studies such as top quark and flavor physics research. The impact of undesired beam-induced background particles on the detector was simulated with a particle gun. Particle trajectories were reconstructed for different amounts of background with the ultimate goal of determining track parameter resolution thresholds to inform future designs and necessary tolerances.

Tbx1 gene regulatory network in mandibular development

Jaidyn M. Woods

Mentors: Marianne Bronner and Sierra S. Marable

Tbx1 is a transcription factor implicated in mandibular development within the first pharyngeal arch. Loss of Tbx1 function is associated with craniofacial and cardiovascular abnormalities, including defects observed in 22q11.2 deletion syndrome, yet its exact role in mandibular development remains unclear. Conflicting studies have reported Tbx1 expression in either the mesodermal core of the pharyngeal arches or subsets of neural crest-derived cells. To characterize spatiotemporal expression of Tbx1 in the developing chick mandible, hybridization chain reaction fluorescence in situ hybridization and Sox10E2-GFP in ovo electroporation followed by whole-mount immunostaining were performed in HH14-HH22 embryos. Tbx1 did not overlap with Sox10-positive cranial neural crest cells in the pharyngeal arches but was detected in GFP-positive neural crest derivatives, suggesting expression after Sox10 downregulation. Ongoing experiments using mesodermal markers will determine whether Tbx1 is also expressed in the mesodermal core. CRISPR inhibition will be used to knock down Tbx1 and evaluate its effects on Sox10 expression, mandibular morphology, and candidate downstream genes to identify the cell populations and gene networks disrupted during mandibular development. Together, this work explores the molecular and cellular functions of Tbx1 during mandibular development and provides insight into how disruption of Tbx1 may contribute to craniofacial defects.

Learning-based methods for ground-effect estimation on thrust-vector-controlled rocket vehicles

William A. Wrolstad

Mentors: Mory Gharib and Jack N. Caldwell

Precise vertical landing of small rocket vehicles requires accurate control in the presence of aerodynamic disturbances. As these vehicles approach the ground, aerodynamic ground-effect forces become a significant source of disturbance, and accounting for these forces in the controller may improve landing performance. This project advances the Daedalus family of electric rocket simulators as a test platform for autonomous precision landing algorithms. During the project, electrical systems for the next-generation Daedalus3 vehicle were developed, and the baseline waypoint controller for the Daedalus2 platform was implemented to support autonomous flight testing. This project also implements learning-based control methods to estimate and account for ground-effect forces on Daedalus2. Previously, similar learning-based methods have been implemented on quadcopters, where the thrust vector is fixed relative to the vehicle orientation, enabling ground-effect force estimation from the vehicle state alone. In contrast, thrust-vector-controlled vehicles decouple thrust direction from vehicle orientation, requiring additional information, such as the thrust vector control

outputs, to accurately predict ground-effect forces. The resulting model will be integrated into the landing controller on Daedalus2 and evaluated for its ability to improve touchdown precision. This work extends learning-based ground-effect estimation to a broader class of vertical-landing vehicles and provides a foundation for future testing on the Daedalus3 platform.

ATLAS ITk pixel testing and laser characterization of the NAPA-p2 monolithic active pixel sensor

Emily J. Wu

Mentors: Charles Young, Harvey B. Newman, and Caterina Vernieri

Sensors for high energy physics experiments must meet benchmarks for sufficient spatial resolution and radiation hardness while maintaining low material and power budget. As the LHC upgrades to HL-LHC, significant increases in pp interactions and instantaneous luminosity will demand detector improvements. ATLAS will replace its inner detector with an all-silicon Inner Tracker (ITk), featuring an inner pixel detector comprised of hybrid pixel modules. Hybrid pixel modules are composed of separate front-end CMOS chips and passive, high-resistivity silicon sensors joined through bump-bonding, with a flexible PCB serving as the module connection interface. ITk will face the highest radiation fluence and pileup in the detector system, necessitating quality control (QC) to ensure performance within specifications. We perform electrical QC tests to evaluate data link quality, functionality of analog and digital electronics, equivalent noise charge levels, sensor dynamic range via leakage current and breakdown voltage, and bump bond connectivity. This characterization provides a systematic approach for validating detector modules.

While ITk employs traditional hybrid pixel modules, development of monolithic active pixel sensors (MAPS) is ongoing. In MAPS, the pixel sensor and front-end readout chip are integrated on a common silicon substrate and can be fabricated using standard CMOS processes, offering benefits in noise, cost, material budget, and power efficiency. Ongoing analysis of the 65nm NAPA-p2 MAPS prototype seeks to characterize pixel efficiency, time walk, and jitter. A 1080nm pulsed laser is delivered through a fiber and focused on a selected pixel to generate signal charge. These measurements will inform optimization of MAPS and future silicon pixel detectors for targeted implementation in e+e- colliders.

Implementation of an observation simulator in ClimaAtmos.jl

Hung-Chih Wu

Mentors: Tapio Schneider and Zhaoyi Shen

Clouds remain a major source of uncertainty in climate projections because many physical processes within clouds occur at scales smaller than a model grid cell. Climate models therefore represent these processes using parameterizations, which must be carefully evaluated. Satellite observations provide broad and consistent measurements of clouds and are therefore valuable for this purpose. However, satellites do not directly measure most geophysical quantities of interest. Instead, they measure electromagnetic signals from which these quantities are inferred using retrieval algorithms, making direct comparisons with model output difficult. Observation simulators address this problem by converting model output into quantities that a satellite instrument would observe. In this project, we implement a simulator for the CloudSat Cloud Profiling Radar (CPR) in ClimaAtmos.jl, the atmospheric model developed by the Climate Modeling Alliance (CliMA). The simulator represents unresolved cloud and precipitation variability, distributes hydrometeors within each model column, and calculates radar reflectivity and attenuation. Comparisons with an established simulator show similar overall reflectivity patterns, with the closest agreement for rain and larger differences for cloud liquid and cloud ice. This implementation enables more direct evaluation of cloud and precipitation parameterizations in ClimaAtmos using CloudSat observations.

The Type-Theory Virtual Machine: A logic-aware, compiler-like system for theorem-prover interoperability

Jiayi Wu

Mentors: Anima Anandkumar and Robert Joseph George

Formal mathematics remains fragmented across theorem-prover ecosystems, and the benchmark ITPEval shows that current language models struggle especially with translations involving ecosystem-specific dependencies. Therefore, we present the **Type-Theory Virtual Machine (TTVM)**, an end-to-end translation pipeline that preserves declaration dependencies and foundation-specific assumptions across provers: starting from exporter and checker evidence, TTVM builds a checked source unit containing a theorem and its required declarations; the Core Intermediate Representation (CoreIR) then normalizes supported logical structure while preserving provenance and keeping unknown symbols opaque; at the same time, AtlasTT evaluates transport between the source and target logical worlds, where effects record logical and trust obligations, morphisms declare transport paths, and residues identify unresolved incompatibilities. Together, the AtlasTT decision and dependency graph determine a target-neutral transport plan. A compatible backend then renders the ordered declarations and theorem in target syntax for checking by the target prover. This layered design replaces pairwise rules with reusable source decoders, foundation-aware transport, and target backends, establishing a scalable basis for cross-ITP statement translation. Overall, TTVM enables formal results to be reused across prover ecosystems without collapsing distinct logical foundations or hiding the assumptions on which they depend.

Reduced-order temporal forecasting of Southern North Sea velocity fields

Justin J. Wu

Mentors: Steven Brunton, Matthew M. Gherman, and Freja Petersen

Ocean velocity fields are high-dimensional, making reconstruction and forecasting computationally expensive. This project investigates whether compressed representations of Southern North Sea velocity fields can reconstruct current flow states and predict future ones. Velocity snapshots were interpolated onto a 64 by 64 grid with two velocity components and represented using Proper Orthogonal Decomposition (POD), which finds dominant spatial patterns, and a convolutional autoencoder, which learns a compressed representation from the gridded data. These methods were first evaluated on reconstruction accuracy. Linear time-update models were then fit in the reduced coordinates and evaluated at forecast horizons of 1, 5, 10, and 20 steps. POD performed best for reconstruction and single-step forecasting, while the convolutional autoencoder achieved lower forecast error across longer horizons. Additional experiments tested reduced external forcing variables and a nonlinear reduced-order model, but these did not consistently improve prediction. These results suggest that representation choice strongly affects reconstruction quality and forecast stability when modeling ocean flow.

Neural network variational Monte Carlo in discrete bases

Qianxi Wu

Mentor: Sandeep Sharma

Neural-network variational Monte Carlo (VMC) replaces the fixed functional forms of traditional trial wavefunctions with a parameterized network. In practice such calculations are limited less by the flexibility of the ansatz than by the cost of the local-energy estimator. Two representations are in common use. Real-space methods parameterize continuous electron coordinates and impose antisymmetry through determinants of configuration-dependent orbitals; their local energy requires a Laplacian, costing $\mathcal{O}(3N_e)$ derivative passes through the network. Occupation-space methods parameterize an amplitude for each discrete configuration and need no derivatives, but a general two-electron Hamiltonian connects each configuration to $\mathcal{O}(N_e^2 N_v^2)$ others, so each local energy costs that many amplitude evaluations.

We surveyed the wavefunctions, samplers, optimizers, and estimators used in both families, and reports measurements on Be, Be_2 . The gradient estimator, not the ansatz, sets the floor, and its degradation is heavy-tailed because $|\Psi|^2$ concentrates on few configurations. The corresponding quantity for the local energy is the cancellation between its large diagonal and connected contributions, which is markedly incomplete for occupation-space networks.

A Gausset basis addresses this cost structure directly. Because Gaussets are localized and approximately orthogonal, the two-electron interaction reduces to density–density terms, so the interaction is diagonal in the occupations and configurations are connected only by one-electron hops. The number of connections is then $\mathcal{O}(N_{\text{es}})$, set by the kinetic stencil and independent of the number of sites, and no Laplacian is required. We develop this into a working method with FermiNet NN backbone, a semistochastic estimator with a Hartree–Fock control variate evaluates the remaining connected sum at order 10^2 network evaluations per walker, under an identity that also permits clipping the reference ratio without introducing bias. Trained with KFAC, the resulting wavefunction follows a continuum FermiNet trajectory on Be and Be_2 .

Sub-grid black holes: Learning the multi-scale dynamics of accretion feedback

Chuyang Xiang

Mentors: Anima Anandkumar, Chuwei Wang, Elias R. Most, and Haiyang Wang

Supermassive black holes regulate galaxy evolution through a coupled feeding–feedback loop that spans from milliparsec-scale horizon physics to kiloparsec-scale galactic environments. Direct end-to-end GRMHD simulations across this range are computationally prohibitive, necessitating subgrid closures at unresolved inner boundaries. Existing methods, such as the cyclic zoom approach, freeze the hydrodynamic boundary state during zoom-out stages, filtering out short-time variability in jet power, orientation, and magnetization that is essential for realistic feedback. We introduce a multi-level neural-operator framework that learns both directions of the scale interface in the native spherical geometry of accretion flows. Two complementary operators are developed: a time-conditioned spherical evolution operator G_θ that advances the GRMHD state on spherical shells and provides a dynamically evolving boundary trajectory to the outer solver, and a boundary-to-interior reconstruction operator F_ϕ that maps an outer boundary snapshot and its conserved flux vector to the full 3D interior state at the same physical time. Both operators leverage spherical-harmonic representations for angular structure and are trained with spectral ℓ_1 losses and jet-region weighted objectives. For G_θ especially, we're introducing pushforward loss and semigroup loss to achieve stable long rollouts. Autoregressive rollouts demonstrate sustained physical consistency over tens of thousands of model calls, preserving jet power-law correlations and phase-diagram distributions. The framework replaces stage-wise frozen boundaries with learned, time-dependent subgrid conditions and provides a learned spatial return map during zoom-in, enabling dynamic feedback coupling across scales.

Andrews-Gordon identities and cylindric partitions

Yingshan Xiao

Mentors: Tony Yue Yu, Weihong Xu, and Yifeng Huang

Rogers-Ramanujan type identities are identities of q -series equating an infinite sum with an infinite product and have a variety of applications in integer partitions and many other fields of mathematics. Among them are Andrews-Gordon identities which are direct generalizations of the original two Rogers-Ramanujan identities, and they are of primary concern to this project. One proof of them relates the entire sum side to a generating function of cylindric partitions which are previously known to be the product side. In our work, we aim to refine this connection by interpreting each individual summand of the Andrews-Gordon identity in terms of a subset of cylindric partitions. Eventually, we hope this can be applied to give a proof of further generalizations of identities due to Huang–Jiang–Oblomkov.

Linearithmic scaling density functional theory via Fourier neural operators

Evan B. Xie

Mentors: Anima Anandkumar and Danish Khan

Kohn–Sham density functional theory (DFT) is limited by the repeated diagonalization of the Kohn–Sham Hamiltonian within the self-consistent field (SCF) loop, a step that scales cubically with system size and, for periodic systems, must be performed independently at every sampled k-point. We introduce a Fourier neural operator (FNO) surrogate that replaces this diagonalization step with a single near-linear operator evaluation, mapping the Kohn–Sham potential directly to the output electron density while leaving the outer SCF loop intact. For periodic systems, the surrogate predicts the already Brillouin-zone-summed density, eliminating the need for explicit k-point summations, achieving even greater speedups in compute time. Across molecular and periodic benchmarks, the resulting FNO-DFT method achieves accurate converged densities, total energies, and band structures at 86-fold reduced computational cost relative to conventional DFT, and outperforms graph neural network baselines trained to predict total energy directly from crystal structure. Our work enables accurate, first-principles electronic structure calculations at system sizes and timescales far beyond the reach of conventional DFT, opening the door to routine quantum-mechanical simulation of large, complex, and previously intractable materials systems.

Breaking the degeneracy: Cosmic variance of galaxy UV luminosity functions at Cosmic Dawn in the LUMINA(-EDE) simulations

Jiahui Xie

Mentors: Lars Hernquist, Philip F. Hopkins, Oliver Zier, and Xuejian (Jacob) Shen

Recent JWST observations have revealed a surprisingly large abundance of UV-luminous galaxies at $z > 10$ than what canonical galaxy formation models in the concordant Λ CDM model predict. Several physically motivated extensions to the galaxy formation model or cosmology have been proposed to explain this tension; however, substantial degeneracies remain. Cosmic variance, field-to-field fluctuations in galaxy counts, provides powerful and independent constraints on the host dark matter halo of UV-luminous galaxies. Using the large-volume (500 cMpc^3) cosmological radiation-hydrodynamic simulation suite LUMINA, we systematically evaluate three models in previous work proposed to address the UV-luminous galaxy tension, including enhanced star-formation efficiency, UV variability from bursty star-formation, and cosmological extensions that enhance small-scale primordial power spectrum, early dark energy. We make predictions for cosmic variance in fields of view matching existing JWST/NIRCam and planned Roman surveys, sampling each field as random-position oriented pencil beams with corrections for field overlaps. We compare our results with predictions from linear perturbation theories and halo bias models, $\sigma_{cv}^{linear} = b_{eff}\sigma_m$, and demonstrate that cosmic variances serve as a complementary probe of the host halo of UV-luminous galaxies to measurements of large-scale galaxy clustering, providing a new route to constrain galaxy formation theories and cosmological models at cosmic dawn.

Integrating visual perception with iterative factor-graph reasoning for puzzle solving and constructing verifiable benchmarks of general reasoning

Ruohan Xie

Mentor: Tony Yue Yu

Visual puzzles require reading images, relating distant clues, and choosing actions whose consequences unfold over many steps. This project develops a shared visual puzzle solver and a factory for verifiable tests of general reasoning. The proof-of-concept solver receives rendered images rather than symbolic boards. A visual encoder reads the whole image and samples features at cell and rule locations, preserving readings as probability distributions. These readings enter a factor graph, a network whose variable nodes represent puzzle quantities and whose factor nodes represent rules. Repeated message passing lets rules revise uncertain visual beliefs. After each edit or move, an exact puzzle engine updates and rerenders the board. Rule inconsistencies select a high-resolution region for the next visual reading. The completed proof of concept validates the effectiveness of this feedback among perception, reasoning, and action across several puzzle families. The Puzzle Factory searches a constrained grammar for compatible rule compositions and constructs solvable instances by beginning

from valid goals and applying verified reverse transitions. Independent replay checks every hidden solution. Its audited collection contains 34 candidate new puzzle types under the project definition of rule-level novelty, represented by 136 formal puzzles testing visual interpretation, rule composition, planning, and global verification.

Making 3D human pose recovery fast enough to work while you move

Xinran Xie

Mentors: Pietro Perona and Ilona Demler

Traditional transformer-based video Human Mesh Recovery (HMR) models are accurate but slow. This makes real-time reconstruction, critical for tasks such as robotics, AR/VR, or sports analysis, an open challenge. Typically, these models have a temporal module in which tokens attend to each other across time. This results in a quadratic runtime cost with the number of frames in a video and a heavy memory demand. We deliver the first attempt to adapt these models into streamable Hybrid State Space Models (SSMs) without compensating on reconstruction quality. Distilled from a teacher model, we fine-tune a fully SSM temporal module on just 8.5k frames and achieve a constant state per tracked person with no key-value cache. We run up to 34 times faster than state-of-the-art HMRS on end-to-end reconstruction tasks without sacrificing performance on standard benchmarks such as RICH and EMDB. The efficiency gain becomes more significant with larger context size, which we show is beneficial for reconstruction quality. Furthermore, we propose not only a single model, but a transferrable recipe for converting an offline teacher HMR model into an efficient and streamable lightweight reconstruction module with preserved quality, minimal re-training, and the flexibility for users to choose the optimal efficiency to reconstruction quality tradeoff.

Mechanistic and kinetic characterization of a vanadium deprotonated-metallacyclobutadiene catalyst for cyclic polyalkyne synthesis

Alex S. Xu

Mentors: Daniel Mindiola, Jonas C. Peters, and Dylan Grigas

Cyclic polyalkynes are a family of conjugated polymers notable for their diverse physical properties, such as superior electrical conductivity and flexibility, presenting potential applications in a variety of materials-related fields. Current catalysts often display uncontrolled kinetic behavior leading to mechanistic intractability or utilize expensive and impractical 4d and 5d metals. In this study, we present the analysis of a vanadium deprotonated-metallacyclobutadiene catalyst that operates via a ring-expansion pathway. The catalyst demonstrates broad alkyne substrate scope, including monomers with aryl halide-, anisole-, polycyclic aryl-, and heterocycle- substituents. NMR analysis and the synthesis of a catalytically active intermediate have allowed for an extensive kinetic study, providing insight into the reaction mechanism. Together, these studies provide a detailed mechanistic understanding of vanadium-catalyzed cyclic polyalkyne synthesis, establishing wide functional group tolerance on an economically efficient metal and allowing for focused optimization of the reaction.

Preparing for the deluge of brown dwarfs in NEO Surveyor data

Allison A. Xu

Mentors: Joseph Kirkpatrick, Christopher Gelino, and Federico Marocco

Understanding star formation is a central goal of modern astronomy, as it shapes stellar and sub-stellar populations throughout the Galaxy. Brown dwarfs, which bridge the gap between giant planets and low-mass stars, remain poorly understood because few have been observed at the lowest masses. The Near-Earth Object Surveyor (NEOS), scheduled to launch in September 2027, offers improved sensitivity to brown dwarfs, making it important to evaluate its expected performance before launch. Our work uses simulated NEOS observations to assess its ability to detect and characterize brown dwarfs. We verified the known brown dwarfs and white dwarfs (included as a proxy for cold brown dwarfs) that made it through the simulator to create and calibrate aperture photometry code such that it correctly measures flux values compared to the known values. Then, we successfully injected

increasingly fainter sources using the NEOS point spread function, and began co-adding images to improve sensitivity for faint brown dwarfs and determine how the signal to noise ratio improves with multiple observations. Further work involves difference imaging and moving co-adds to better recover faint sources once background objects have been removed.

Tracial joint spectral measure for nearly commuting matrices

Chenning Xu

Mentor: Otte Heinävaara

Tracial joint spectral measures are measures on the plane and originally introduced by Heinävaara to prove certain trace inequalities for two matrices. Heinävaara gave complicated formulas for the density and singular part of the measure with respect to Lebesgue measure; both formulas involve the eigenvalues of an intermediate matrix. We study the density when the matrices nearly commute. We find a low dimensional approximation of the density and prove geometric properties of the support of density. Specifically, we prove the support is located near certain segments up to first order. Up to second order, the support resembles a union of ellipses, and we find formulas for the ellipses.

Studying the quantum kinetic evolution of neutrinos

Olivia J. Xu

Mentors: Chad Kishimoto and Ryan B. Patterson

In a hot and dense neutrino gas, the evolution of neutrino flavor is described by the coherent growth of quantum phase and a simultaneous scattering-induced destruction of phase. This quantum kinetic evolution is driven by neutrino-neutrino and neutrino-antineutrino interactions, producing non-linear phenomena that could reflect neutrino flavor conversion in the early universe. Using self-consistent solutions to the Quantum Kinetic Equation—a quantum master equation for the evolution of a distribution of neutrinos, we analyze the evolution of initially out-of-equilibrium neutrino distributions to discuss the synchronized flavor oscillations in a purely coherent solution, the path to equilibrium in a purely kinetic solution, and how the quantum kinetic evolution of the neutrinos relates to these behaviors and emergent behavior from the interplay of both. In this project, we aim to describe the interplay between the coherent quantum phase and the scattering-induced decoherence. By exploring the evolution of out-of-equilibrium neutrino distributions, we hope to describe emergent phenomena like those found in the early universe.

Recurrence and transience of excited random walks on trees and lamplighter groups

Owen Xu

Mentors: Matthew Gherman and Tasmin Chu

Multi-excited random walks on a graph G are random walks in which the transition probabilities at any particular vertex x in G depend on the number of times x has been previously visited. A (usually finite) number of cookies $\omega(x, i)$ are placed on each vertex x of G , each with their own specified transition probabilities $\omega(x, e, i)$. At the i th visit of the vertex x , the random walker moves according to the transition probabilities $\omega(x, e, i)$ and consumes the cookie shortly afterward. After all cookies are consumed at a vertex, the walker moves from that vertex with transition probabilities equivalent to that of a simple random walk.

In this project, we examine multi-excited random walks on trees of exponential growth; in particular, we show that on a tree of exponential growth, there exists a recurrent multi-excited random walk with finitely many cookies on each vertex. Separately, we also look at the once-excited random walk on the lamplighter group $\mathbb{Z}/2\mathbb{Z} \wr \mathbb{Z}$ — we provide a streamlined proof of the transience of the simple random walk on the lamplighter group, and we modify this proof to show transience of the once-excited random walk on the lamplighter group.

Optimizing sequential resin–latex extractions from guayule through analytical characterizations

Ruoyun Xu

Mentors: Julia A. Kornfield and Florence Müller

Guayule is a promising natural rubber source alternative. Compared to *Hevea brasiliensis*, guayule is domestic to the US, less susceptible to diseases and its rubber has better mechanical properties and is hypoallergenic. The guayule plant produces two high value products: latex and resin; which both need to be extracted in the highest possible purity for usability and high performance. This work addresses the co-extraction of both products located in the bark through a sequential resin-then-latex extraction method. The aim is to reduce cross-contamination of resin in latex while preserving the resin.

Fresh guayule bark was first sliced, opening resin channels while keeping the latex encapsulated in the parenchymal cells. The bark slices were then submerged in an extraction buffer and immiscible solvent was added on top. The mixture was then agitated to extract the resin while keeping the bark hydrated to dissolve the resin into the solvent. Latex was then extracted through conventional methods and the resin was dried from the solution. This work gives a quantitative comparison of the variation of different process parameters, including the ratio of buffer to solvent and the agitation time. The purity of the latex and resin were analyzed using thermogravimetric analysis and thin layer chromatography, respectively.

By exploring the different process parameters of the sequential extraction process, this work informs important variables for a scale-up of guayule latex and resin extraction.

What if we trained video models on Mars?

Charlie Q. Yang

Mentors: Pietro Perona, Raphaela H. Kang, and Sevan Brodjian

Recent work suggests that foundational video models can develop an intuitive understanding of physical laws. Building upon this, we investigate the capacity of video models to internalize fundamental physical laws by training them in simulated environments with altered physics. To establish a baseline for evaluation, we introduce methods to quantify a model's intuitive physics understanding, including linear probing of Transformer layers to detect impossible physical trajectories and a representation prediction approach (VJEPA-H). We generated synthetic datasets representing universes with distinct physical properties, such as inverted gravity, zero friction, and collisionless environments, to train and evaluate models like V-JEPA and VideoMAE.

Integration of PL spectroscopy and twisted light-assisted STM for nanoscale exciton research on ML-MoS₂

Cheng-Hsun Yang

Mentors: Nai-Chang Yeh and Jen-Te Chang

Monolayer transition metal dichalcogenides (TMDs) host strongly bound excitons and provide a versatile platform for manipulating photoexcited quasiparticles through nanoscale strain and structured optical fields. Nanoscale strain can reshape the local electronic band structure and create energetically favorable regions for quasiparticle accumulation, while light carrying orbital angular momentum provides an additional degree of freedom for controlling excitonic states. Directly correlating these effects requires simultaneous access to the local electronic structure and optical emission, motivating the integration of photoluminescence (PL) spectroscopy with light-assisted scanning tunneling microscopy and spectroscopy (STM/STS).

Here we investigate the experimental and computational foundation for a light-assisted STM-PL platform for monolayer MoS₂ (ML-MoS₂). Gold-assisted transfer is investigated for preparing large-area ML-MoS₂. Finite difference time domain (FDTD) simulations implemented in Meep are used to examine how the tip geometry, tip sample gap, Au termination, excitation wavelength, and spacer structure affect the local electromagnetic field within the STM junction. Complementary time dependent density functional theory (TDDFT) calculations implemented in Octopus are developed to investigate the microscopic optical response of the junction under a maintained bias.

The computational results provide reference for selecting junction geometries and spacer structures that preserve optical field confinement, while the transfer studies identify practical requirements for preparing large and sufficiently clean MoS₂ regions. Together, these efforts establish the sample preparation strategies and modeling framework required for future integration of PL spectroscopy with light-assisted STM and for correlating nanoscale optical emission with the local electronic response of ML-MoS₂.

Speaking to the audience: Mode of address as a partisan signal in U.S. political podcasts

Danielle S. Yang

Mentor: R. Michael Alvarez

Political podcasts are a major but rarely studied news source. We investigate how hosts differ in the way they speak to listeners across the political spectrum. From transcripts of 204 U.S. political podcasts (2018–2026), we measure two registers using published dictionaries and syntactic parsing: hostility and listener-directed address (the second-person instruction aimed at the audience). Partisanship is validated against hosts' campaign contributions and listeners' reported party. Linking the corpus to two national surveys ($n \approx 18,800$, $n \approx 9,000$), we find that conservatives and liberals differ not in hostility but in listener-directed address ($d = 0.78$), a stable trait of each show. Podcast listeners are 0.22 SD more cynical about institutions than other news consumers, in both parties and at every education level. Within political podcasts, listener-directed address predicts a less educated audience (-0.27 SD per SD, income held constant) and, exploratorily, listeners who distrust elections and institutions ($+0.27$ and $+0.22$ SD per SD). Thus, podcasting sorts listeners by how hosts speak, not only what they argue. Because address style is measurable at scale and stable within shows, we show that political register offers a tractable place to study how media environments and institutional distrust are linked.

Wavelength-tunable near-infrared femtosecond stimulated Raman spectroscopy of LiNbO₃

Xinyi Yang

Mentors: David Hsieh and Mingyao Guo

Collective excitations play an important role in quantum materials. Femtosecond stimulated Raman spectroscopy (FSRS) provides a sensitive way to measure these modes by combining a narrowband pump pulse to stimulate vibrational transitions and a broadband femtosecond probe pulse to detect the resulting gain and loss spectrum. Compared with conventional spontaneous Raman scattering, this stimulated process produces a stronger, directional signal while retaining simultaneous spectral and temporal resolution to resolve their dynamics out of equilibrium. In this project, I developed a broadband wavelength-tunable near-infrared implementation of this technique for LiNbO₃. The broadband probe is produced through supercontinuum generation in an yttrium aluminum garnet (YAG) crystal. Raman-pump wavelengths from approximately 1250 to 1500 nanometers will be utilized, and pump-on and pump-off probe spectra will be compared to extract phonon frequencies, linewidths, and intensities. Measuring the Raman response across this tuning range will reveal how the line shapes of different phonon modes depend on excitation wavelength. These measurements will establish the performance of the tunable spectrometer and pave the way for time-resolving the phonon dynamics in LiNbO₃.

Fully-coupled oscillations of perturbed Kerr-Newman black holes

Yi Wei Yang

Mentors: Yanbei Chen and Colin M. Weller

Perturbed black holes undergo a characteristic ringdown process to return to their equilibrium state, emitting a discrete spectrum of quasinormal modes (QNMs) dependent on the geometry of their spacetime. We investigate the free gravito-electromagnetic response of Kerr-Newman (KN) black holes, which can be described by a set of two coupled gauge-invariant partial differential equations. These equations lack the separability available in the Kerr and Reissner-Nordström (RN) limits. We construct a two-dimensional Chebyshev pseudospectral solver that imposes ingoing-horizon, outgoing-asymptotic, and regular-polar boundary conditions. QNM frequencies are obtained by minimizing the relative smallest singular value of the discretized coupled operator, while the associated singular vectors provide the mode eigenfunctions. The solver is validated using convergence and boundary-residual tests, and by recovering established Kerr and RN spectra, as well as published KN mode sequences. Using the coupled eigenfunctions, we reconstruct the gravitational and electromagnetic radiation variables Ψ_2 and Ψ_{-1} , generate free ringdown waveforms and angular emission patterns, and examine how spin and charge can modify the continuity of the spectrum of KN QNMs. This framework provides a route towards a KN Green's function and, ultimately, sourced waveforms for compact objects interacting with charged, rotating black holes.

Energetics of gradients in biology

Zitao Yang

Mentors: Rob B. Phillips, Minakshi Ashok, and Sara Mahdavi

Energy investment is necessary for biological systems to exhibit beautiful and complex patterns. Many of such patterns play crucial roles in the functions of living organisms. For instance, the Bicoid protein gradient in *Drosophila* embryo is crucial in setting up the anterior-posterior axis. Much previous work has been done to understand the molecular dynamics that give rise to the gradients. However, little is known about the power needed to build and sustain them. In this work, we first lay down the theoretical framework to compute the rate of change for various thermodynamic quantities. Specifically, we compare the minimum power needed to maintain the gradient and actual chemical work done on the system. We apply this framework to analyze several biological systems including: Bicoid protein gradient in developing *Drosophila* embryo, kinesin-2 gradient in *Chlamydomonas* flagella, polarity protein Cdc42 gradient in yeast, and Min protein gradient in *E. coli*. A careful account of energy budget reveals an apparent inefficiency of biological approaches to maintaining gradients.

Predicting eye movements from personality traits in autism

Jovytta E. Yauwanta

Mentors: Ralph Adolphs and Qianying Wu

Atypical gaze patterns are widely observed in individuals with autism spectrum disorder, making them a promising biomarker. However, previous studies often overlook the role of personality traits - another important source of individual differences - in shaping gaze behavior. To address these gaps, we analyzed a large-scale webcam-based eye tracking dataset including autistic and non-autistic participants observing group social interaction in the form of multi-speaker zoom conversations. Specifically, we investigated how dimensional personality traits (using Neuroticism Extraversion Openness Five-Factor Inventory) and autistic traits (using Comprehensive Autistic Trait Inventory) in both autistic and non-autistic population mediate visual attention. Current results show that autistic and non-autistic cohorts differ significantly across all five personality dimensions, with the Autistic group exhibiting elevated Neuroticism alongside reduced Conscientiousness, Extraversion, and Agreeableness. Mediation analyses further revealed that Conscientiousness significantly and partially mediated the relationship between autism group and visual attention to distractor stimuli. However, Extraversion was not a significant mediator in the relationship between autism group and visual attention to speaker stimuli.

Optimization of PECVD-grown graphene for 2D heterostructure integration

Tzu-Ning Yen

Mentors: Nai-Chang Yeh and Jen-Te Chang

Graphene-based heterostructures provide a promising platform for investigating the electronic properties of two-dimensional materials. However, their fabrication requires careful control of graphene quality, substrate interactions, and interfacial contamination. This study investigates the growth of graphene on metal substrates using plasma-enhanced chemical vapor deposition (PECVD) as a foundation for the future fabrication of graphene/TMD heterostructures and further applications in strain engineering.

Initial experiments focused on establishing graphene growth conditions on Ni thin films, but growth reproducibility was limited by substrate conditions and instability in the PECVD vacuum system. After identifying and repairing vacuum leaks and restoring stable vacuum conditions, graphene-related D, G, and 2D Raman features were observed more consistently. Subsequent experiments systematically investigated the effects of CH₄/H₂ gas composition while maintaining fixed plasma power and growth time. A more quantitative characterization approach was also adopted by considering Raman peak intensity ratios, peak widths, and spatial uniformity across the samples. The current results establish a foundation for the systematic optimization of PECVD graphene growth conditions and for evaluating direct-growth and transfer-based fabrication approaches toward graphene/TMD heterostructures on different substrate structures.

First-principles modeling of phonon-mediated anomalous Hall transport in magnetic materials

Chieh-An Yu

Mentors: Marco Bernardi and Lauren A. Tan

The anomalous Hall effect (AHE) is a fundamental transport phenomenon in magnetic materials, where an applied electric field generates a transverse voltage through the interplay between spin-orbit coupling, electronic band structure, and carrier scattering. While intrinsic AHE arising from Berry curvature and conventional extrinsic mechanisms caused by impurity scattering have been extensively studied, the influence of phonon-mediated scattering on anomalous Hall transport at finite temperature remains an open theoretical challenge.

In this project, we investigate anomalous Hall transport from both theoretical and first-principles perspectives. We develop a theoretical framework for phonon-mediated extrinsic AHE by deriving the contributions from skew scattering and side-jump mechanisms within the semiclassical Boltzmann transport formalism. This analysis highlights the role of inelastic electron-phonon interactions and higher-order quantum interference processes in generating asymmetric carrier scattering.

In parallel, we perform first-principles calculations of magnetic materials using Quantum ESPRESSO, Wannier90, and Perturbo. We calculate electronic band structures, Berry curvature distributions, and intrinsic anomalous Hall conductivity for ferromagnetic bcc iron and the kagome ferromagnet Fe₃Sn₂, a material known for its strong spin-orbit coupling and large anomalous Hall response. These calculations provide insight into the electronic origins of anomalous Hall transport and establish a computational foundation for future investigations of phonon-mediated extrinsic contributions.

Persona- and urgency-driven sycophancy in frontier LLMs

Clara Yu

Mentors: Ankit Sakhuja, Shinsuke Shimojo, and Nimay Hazare

The integration of Large Language Models (LLMs) into clinical workflows introduces sophisticated capabilities for diagnostic reasoning and patient communication, yet a critical vulnerability lies in the inconsistency of their internal knowledge states across varying linguistic frames. This paper investigates persona- and urgency-driven sycophancy in frontier LLMs, evaluating over 48,000 responses to examine how role-playing constraints and perceived situational urgency affect core clinical accuracy. Our results show that 12.0% of all responses displayed sycophantic behavior, with

susceptibility varying by more than an order of magnitude across models, ranging from 1.4% sycophancy in Claude 4.6 Sonnet to 38.4% in Gemini 3.5 Flash. Furthermore, prompting models to interact with high-status clinical personas significantly increased sycophantic behavior compared to lay-person personas, peaking with the Senior Physician persona. This effect increased monotonically with situational urgency, from a 4.3% sycophancy rate at Level 0 to 23.8% at Level 4 (OR 18.90), with critical clinical reversals accounting for 63% of failures. Ultimately, these findings indicate that stylistic prompting significantly degrades the consistency of medical reasoning, highlighting the urgent need for future alignment strategies to safeguard core knowledge states before deploying LLMs as clinical decision support systems.

Why do buyers discount flood-resistant homes? A housing choice experiment

Emily Q. Yu

Mentor: Hannah Druckenmiller

Home elevation is among the most effective physical adaptations for reducing flood damage, yet elevated homes sell at a discount. This project uses a survey choice experiment to investigate how prospective home buyers evaluate residential properties with flood-resistant construction. We recruit U.S. adults aged 25 or older who purchased or seriously considered purchasing a home in the past five years. Participants complete six choice tasks, each presenting a pair of hypothetical home listings in which visible elevation is varied through listing photographs, alongside information on monthly flood insurance premiums and other standard housing attributes. Primary outcomes are the home selected from each pair and a 0–100 desirability rating for each home. We test whether visible elevation reduces the probability that a home is selected and lowers its rating, and whether these effects vary with household mobility limitations, prior flood damage experience, and risk aversion. The findings will clarify the role of nonfinancial concerns, including appearance and accessibility, in explaining why buyers discount elevated homes despite their flood-protection benefits.

Extending temperature control capabilities for Laser-driven Ultrafast Impedance Spectroscopy

Jack S. Yu

Mentors: Scott K. Cushing, Amy Lin, and Natan Spear

Solid-state electrolytes are promising materials for next-generation batteries due to their potential to improve safety and energy density compared to conventional liquid electrolytes. However, the microscopic mechanisms that enable rapid ion transport through these materials remain incompletely understood. Recent studies suggest that lattice vibrations, or phonons, influence how ions move through a crystal structure, but directly determining how these vibrations affect ionic transport remains experimentally challenging. This project develops a temperature-controlled sample stage for Laser-driven Ultrafast Impedance Spectroscopy (LUIS), a technique that uses laser excitation to probe changes in ion hopping on ultrafast timescales. The stage uses thermoelectric cooling to control the temperature of solid-state electrolyte samples, allowing the equilibrium population of thermal phonons to be systematically varied. Temperature-dependent LUIS measurements will be demonstrated using systematic selenium substitutions in sodium thiophosphate ($\text{Na}_3\text{PS}_{4-x}\text{Se}_x$), a highly conductive solid-state electrolyte, to examine how changes in the initial phonon population influence phonon-ion coupling. Ultimately, the integration of temperature control into LUIS provides an additional experimental parameter to systematically compare ion transport across different temperatures and materials, enabling a better understanding of the factors governing ionic conduction.

Identifying and steering with interpretable features in protein models using sparse autoencoders

Ryan D. Yu

Mentors: Grant Rotskoff, Kara D. Fong, Jason Liu, and Soojung Yang

Generative protein models have recently become one of the most important advances in modern biology. In particular, AlphaFold has become incredibly important for predicting protein structures without lengthy spectroscopy experiments. Other models like BioEmu add additional functionality like generating protein samples along a Boltzmann distribution. However, these models are still mostly

black boxes. To improve general understanding of these models, we train sparse autoencoders on these generative protein models to identify features corresponding to structural and thermodynamic properties, as well as at what point in the model these features appear. From this, we show that these features can be used to steer generation towards desired properties.

TMED8 selectively promotes kappa opioid receptor biogenesis

Zhenxi Yuan

Mentors: Rebecca M. Voorhees and Aline Milach Teixeira

Kappa opioid receptor OPRK1 (KOR) is a Class A G-protein-coupled receptor (GPCR) involved in mood, anxiety, and addiction. While much attention has focused on how KOR responds to endogenous ligands and drugs, less is known about how cells regulate whether newly synthesized KORs reach the cell surface in the first place. Understanding the biogenesis factors that control KOR abundance could reveal how cells tune receptor availability before signaling and may eventually suggest new strategies for modulating opioid receptor activity in disease. Through a CRISPRi screen for regulators of KOR total abundance in K562 cells, we identified TMED8 as a candidate positive regulator of KOR biogenesis. TMED8 is poorly characterized and atypical relative to canonical p24 family proteins, making its cellular functions unclear. Initial TMED8 knockdown and overexpression experiments both suggest that TMED8 positively regulates KOR total abundance, with a more selective effect on KOR than on other tested opioid receptors and GPCRs. Further preliminary imaging indicates that TMED8 localizes near early secretory pathway compartments, including the ER and ER-Golgi intermediate compartment. These findings support a model where TMED8 promotes KOR biogenesis, potentially by acting within the early secretory pathway. Ongoing work will test this working model using perturbation and trafficking assays to define where and how TMED8 participates in KOR biogenesis.

A systematic evaluation of hardware and acquisition parameters for near-infrared imaging

Vivian C. Yue

Mentors: Anuradha Godavarty and Lihong Wang

Near-infrared (NIR) imaging allows light to penetrate several millimeters to centimeters into tissue due to reduced absorption and scattering in the 650–1000 nm spectrum, which allows for noninvasive assessment of subsurface physiological information. The objective of this project is to optimize near-infrared (NIR) imaging by identifying the imaging hardware, acquisition parameters, and illumination methods that produce the most consistent high quality images. Camera systems were evaluated under white light and multiple NIR wavelengths (660, 690, 830, and 850 nm) by comparing image quality using different image formats, shutter modes, and downsampling methods during both static and continuous imaging. The optimal camera settings and the selection of cameras that provide the best performance for NIR imaging have been identified through this process. Additional work included evaluating optical filters, alternative LED light sources, compact white-light cameras, and time-of-flight sensors, while developing synchronized imaging software capable of controlling multiplexed illumination and image acquisition. Custom illumination boards and integrated control electronics have also been developed and are awaiting experimental validation. Future work will focus on achieving stability in the illumination with replacement light sources, and integrating the optimized hardware and software into a standardized imaging protocol that provides clean and efficient near infrared imaging results.

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.

The funding crisis in human embryonic stem cell research: Historical patterns and the path forward

Simon D. Zernicki-Glover

Mentor: David Zierler

Human embryonic stem cell (hESC) research provides critical insight into early human development, pregnancy loss, congenital disease, drug response, and regenerative medicine. However, the field is vulnerable to political restrictions, unstable federal support, public misunderstanding, and pressure to demonstrate immediate clinical applications. This project examines how the historical development of American science funding created these vulnerabilities and identifies ways to protect important research that existing funding systems aren't covering. Drawing on historical and policy research, federal budget data, and interviews with developmental biologist Marianne Bronner and Caltech development officer Laura Yoo, I compare the roles of federal funding, state programs, industry, and philanthropy. I find that no single funding source can support the entire hESC research pipeline. Federal funding provides scale, legitimacy, and peer review but remains legally restricted. State programs can create regional alternatives, while industry generally favors research with clearer commercial potential. Philanthropy cannot replace federal science funding, but it can provide flexible support for politically vulnerable and long-horizon research. I propose a coordinated funding ecosystem in which philanthropic organisations prioritise current-use funding, long-term endowments, early-career researchers, donor education, and shared scientific infrastructure. Such support can make vulnerable science fundable before its full value becomes obvious.

Thermal imaging with a temperature-sensitive polymer using electrical impedance tomography

Tara T. Zhan

Mentors: Chiara Daraio and Brian Ahn

Temperature mapping across large, flexible surfaces would benefit robotics and wearable sensors for healthcare, but infrared cameras cannot image curved or obstructed objects, and grids of point sensors require complex wiring that scales with area. We are developing a thermal imaging sheet that avoids both limits by pairing electrical impedance tomography with a temperature-sensitive polymer whose conductivity increases as it is heated. A thin film of the polymer is cast onto a flexible board with sixteen perimeter electrodes. Heating raises the conductivity only where the heat lands, altering how current spreads between every pair of electrodes, so an image of the interior can be reconstructed from edge measurements. A known heat source is applied at chosen positions while impedance is recorded across all electrode pairings. Reconstructions place a single localized region of raised conductivity at the heat source, and its magnitude grows with the temperature applied, showing that the sheet recovers both where heat is applied and how intense it is. Ongoing work quantifies temperature sensitivity and spatial resolution, improves film durability, and extends testing to curved surfaces.

WordBench: A benchmark for learning and search in algebraic word problems

Caroline Zhang

Mentors: Anima Anandkumar and Robert Joseph George

Mathematical discovery systems increasingly combine large language models (LLMs), reinforcement learning (RL), and formal verifiers, yet most mathematical benchmarks remain static collections of problems. We introduce WordBench, a benchmark for algebraic word problems that combines procedurally generated datasets with executable search environments and exact verification. A word problem asks whether finite symbolic expressions, or words, represent the same mathematical object under specified generators and relations. An open example is the Andrews-Curtis conjecture, where a successful certificate is a sequence of allowed presentation moves transforming a relator tuple to the trivial presentation. WordBench contains 24 families spanning rewrite and witness-construction problems, each with generative datasets, a deterministic symbolic search environment, and exact verification of submitted traces, action sequences, or algebraic objects in the Lean theorem prover. The released corpus, WordBench-Lit, contains 13,324 instances, and the underlying data generators can produce fresh instances beyond this set. We benchmark classical breadth-first search, PPO agents, and LLMs under common verifier-defined success criteria. The results reveal strongly solver-dependent empirical difficulty: for example, PPO solves 99.1% of instances on Thompson's group F while some open-source LLMs remain below 10%, whereas frontier LLMs achieve near-perfect performance on Artin-Tits and braid rewriting, where PPO remains below 1%. More broadly, frontier LLMs solve a large fraction of WordBench-Lit, but performance remains nonuniform across families, with Andrews-Curtis being substantially less solved. Beyond benchmark performance, WordBench's environments enable controlled experiments on search behavior, learning dynamics, and empirical relationships across mathematical problems. As a case study, we find that cross-family PPO transfer learning can solve new dataset instances not discovered by target-only training, with different source families providing complementary coverage. Together, these results position WordBench as a common verifier-facing benchmark for studying learning and search across diverse mathematical word problems.

An engineered prodrug-activating enzyme for circuit-based cancer therapy

Evan Z. Zhang

Mentors: Michael B. Elowitz, Andrew Lu, and Lukas Moeller

Enzyme-prodrug therapies, such as yeast cytosine deaminase (yCD) paired with the prodrug 5-fluorocytosine, rely on enzyme-based conversion of a non-toxic prodrug into a potent, diffusible chemotherapeutic within a tumor site to enable greater tumor penetration and clearance. However, existing strategies for restricting enzyme activity to cancer cells, such as tumor-specific promoters or viral delivery, provide limited specificity and cannot directly sense oncogenic signaling within a cell. Therapeutic circuits, which use engineered proteins to detect oncogenic signals and conditionally trigger cell death, offer a potential route to directly couple prodrug activation to the cancer-specific signal. Here, we present the foundation for a mutant Ras-activatable yeast cytosine deaminase, in which deaminase activity is engineered to depend on a sensor module that is specifically activated by oncogenic Ras. We explore several different strategies for engineering circuit-controlled yCD and evaluate dynamic range using a HEK293-based assay. Overall, these results lay the groundwork for a circuit-based enzyme-prodrug therapy that could combine the tumor-killing efficiency of 5-fluorocytosine conversion with the specificity advantages of Ras-sensing protein circuits.

Timing information for background rejection in a muon missing-momentum search

Jinbo Zhang

Mentors: Bertrand Echenard and David G. Hitlin

A muon-beam missing-momentum search for an invisibly decaying dark photon A' must reject reducible backgrounds at the 10^{-7} level. We study what detector information that rejection requires in an idealized Geant4 simulation of a 5 GeV/c μ^- beam on an active LYSO target read out with electromagnetic, downstream hadronic, and side hadronic calorimeters. A graph neural network over a 625-node representation of the four subdetectors, reading energy and position only, leaves a mean of 5.6 background events at 90% matched fired-signal efficiency in a 5.0×10^7 -muon benchmark, with strong seed-to-seed variation. That residual is entirely muon-nuclear: 12 unique events, every one

carrying a neutron at the target exit. Adding two per-node timing observables (the energy-weighted mean deposit delay and its temporal RMS) to an otherwise identical network lowers the residual in four of five seeds and ties in the fifth; the mean falls from 5.6 to 1.6 events, and ten of the twelve disappear entirely. Under Gaussian smearing of the deposit times from 100 ps to 50 ns the mean residual stays between 1.2 and 2.2 throughout, so the effect does not depend on exact simulation timestamps. Because a real channel does not resolve individual deposits, this demonstrates robustness of the temporal information, not a hardware requirement; a pulse-level detector-response model and hadronic-model validation remain necessary before any detector-performance reading.

Disturbance-compensated reinforcement learning control for trajectory tracking

Junze Zhang

Mentor: Soon-Jo Chung

Unmodeled forces, including actuator mismatches, environmental disturbances, and surface resistance, can often amplify a controller's trajectory tracking errors during experimental operation. This project thus develops a novel reinforcement learning-based adaptive control algorithm that improves trajectory tracking by learning and correcting for these differences. The method consists largely of a model-based tracking controller and adds a learning-based disturbance correction. It first uses reinforcement learning (RL) to learn a shared disturbance representation offline from randomized trajectories and continuously varying force and moment disturbances, which improves upon the supervised learning models by learning compensation that accounts for actuator limits and delayed effects without requiring exact disturbance models and expert demonstrations. During operation, a small set of coefficients is updated online to adapt this representation to the current disturbance. The method was then developed and evaluated in simulation before being tested on the Caltech Multi-Spacecraft Testbed for Autonomy Research (M-STAR). Performance was compared to both model-based methods, such as pure PD controller, and learning-based methods, such as pure RL controller and controller using RL to learn the disturbance term, on 1,000 unseen trajectories. The addition of an online coefficient adaptation produced a further modest improvement in simulation. Learning only a disturbance correction improved all four tracking measurements, while using RL to learn the total control command directly produced the largest errors. Initial hardware testing demonstrated an average position error of approximately 0.0698 meters and an average heading error near zero (0.0301 radians) while following a closed trajectory, for both constant and varying desired headings. These results indicate that reinforcement learning is more effective as a correction to an existing controller than as its complete replacement, and that combining offline learning with online adaptation is the most promising approach for further hardware testing and eventual multi-spacecraft fault recovery.

Investigating sex-specific differences in the squid-*Vibrio* symbiosis using 5'-RNA sequencing

Kaifeng Zhang

Mentors: Margaret J. McFall-Ngai and Drew Honson

The symbiotic relationship between the Hawaiian bobtail squid and the luminous bacterium *Vibrio fischeri* provides a well-established model for studying host-microbe interactions. Although this association has been extensively characterized, whether males and females exhibit distinct molecular responses to symbiosis remains largely unknown. The goal of this project is to investigate sex-specific transcriptional differences during symbiosis by combining molecular sex determination with transcriptome analysis. First, genomic DNA was extracted from hatchling squid and quantitative PCR was used to determine the sex of each individual. RNA was then isolated from light organs of squids with different symbiotic states and prepared for 5' RNA sequencing. Sequencing data were analyzed to compare gene expression patterns between males and females and to identify sex-specific responses to bacterial colonization. These findings provide a foundation for understanding how sex influences the establishment and maintenance of beneficial host-microbe interactions.

Tracing accretion-state evolution in the 2021 outburst of GX 339-4 with NICER

Yuchi Zhang

Mentors: Fiona A. Harrison and Shina Adegoke

The evolution of accretion disks and coronae through black hole X-ray binary state transitions remains poorly understood. We investigate the 2021 outburst of GX 339-4 using 262 observations from the Neutron star Interior Composition Explorer (NICER), totaling 680 ks. We model the flow evolution with an absorbed Comptonized multi-temperature disk model. In the soft state, the fitted disk temperature varies between ~ 0.83 keV and ~ 0.46 keV. In this state, we found that the disk luminosity vs temperature relation follows the prediction for a standard thin disk. The intermediate-state observations appear to follow distinct tracks in the disk luminosity vs temperature relation, depending on whether the source was transitioning from the hard to the soft state or from the soft back to the hard state. In the hard state, the cool disk is weakly constrained, and no similarly tight luminosity-temperature relation emerges.

Towards a convergent fragment-coupling approach of grayanane diterpenoids via selective lithium-halogen exchange of a vinyl dibromide

Zhehao Zhang

Mentors: Sarah E. Reisman and Joshua Belfield

Grayananes are a class of natural products (diterpenoids) with a fused 5/7/6/5 carbon skeletal structure, with the central 7-membered ring being of great synthetic challenge. Herein, we report our progress towards a fragment coupling strategy to forge the central 7-membered ring, via selective lithium-halogen exchange into a vinyl dibromide fragment. We report the successful synthesis and lithiation of the proposed vinyl dibromide fragment, as well as progress continuing from last SURF on developing the bicyclo[3.2.1]octane fragment as a viable coupling partner. We envision that with both fragments in hand, two selective couplings will enable access to the Grayanane skeleton, which can be elaborated to targets such as Pieristoxin S.

Langmuir probe for Caltech Water-Ice Dusty Plasma Experiment

Annie W. Zhao

Mentors: Paul M. Bellan and Jessica C. Gonzalez

Langmuir probe measurements of electron temperature (T_e) are being conducted for the Caltech Water-Ice Dusty Plasma Experiment, where ice grains are formed by injecting water vapor into a cold, capacitively coupled plasma (CCP). To better understand the impact of ice grain formation on the surrounding plasma, measurements of T_e were taken with and without the presence of ice grains. Based on preliminary results, a probe bias exceeding +20 V DC is necessary to characterize the plasma's full range of electron energies. Langmuir probes are difficult to implement in CCPs due to I-V curve smearing caused by sheath oscillations¹ and spurious RF pickup by probe electronics. Two methods are being used to reduce this effect. The first method aims to decrease smearing by taking measurements at a fixed phase of the RF cycle over a range of probe biases. The second method attempts to avoid RF effects via data collection in the plasma afterglow. Initial measurements using traditional techniques and the first method both show a higher T_e than expected. Both reduction methods for RF effects will be attempted, refined, and presented.

¹A. Garscadden and K. G. Emeleus, Proceedings of the Physical Society 79, 535 (1962).

Capturing membrane dynamics and spike timing of human neurons at scale using neural operators

Cindy Y. Zhao

Mentors: Anima Anandkumar and Valentin Duruisseaux

Characterizing the properties of neuronal cell types is central to understanding their functional roles in the brain. Yet, scaling biophysically realistic neuron models to large populations while preserving rich cellular dynamics and within-cell-type variability remains a major challenge. We present a cell-type-specific neural operator framework for scalable neuronal population modeling that learns a mapping

between interpretable electrophysiological features and somatic voltage responses elicited by current stimuli. The model generates large ensembles of these responses while preserving variability across cells and model configurations, achieving a $3,640\times$ speedup over numerical solvers. We evaluate the approach using metrics that describe membrane dynamics, action potential waveforms, spike timing accuracy, and input-output firing relationships. The predicted electrophysiological feature distributions show strong agreement with PDE-simulated outputs and human experimental recordings, including high explained variability in properties such as inter-spike intervals, firing frequency, action potential shape, and post-spike recovery. Through efficient ensemble simulation of cell-type-specific neuronal responses, the framework provides a scalable tool for studying cellular diversity, probing variability near functional thresholds, and generating synthetic yet biophysically grounded neuronal populations.

Joint optimization of k-space sampling and water-fat separation for accelerated dual-echo MRI

Xiaoli Zhao

Mentors: Shreyas S. Vasanawala, Houman Owhadi, Aiqi Sun, and Yan Wu

3D dual-echo spoiled gradient echo MRI provides strong T1 contrast and enables water-fat separation, but 3D acquisitions are time-consuming, and conventional Dixon-based methods are susceptible to water-fat swaps. No prior work has built self- and cross-attention between echoes directly into the reconstruction network, letting the two echoes share anatomical information during reconstruction itself. We formulate dual-echo reconstruction as a data-driven inverse problem, benchmarking U-Net and T-Net backbones augmented with self-attention, cross-attention, and a physics-informed fusion gate conditioning attention weights on inter-echo phase difference, trained with data-consistency at $R=10$. T-Net with cross-attention performs best (Echo 0: 29.08 dB PSNR, 0.9512 SSIM, 0.0577 NMSE; Echo 1: 32.63 dB, 0.9856, 0.0168), outperforming the U-Net baseline (20.50 dB, 0.7439, 0.4325); cross-attention alone matches the fusion-gated variant, suggesting inter-echo information sharing is the primary driver of improvement. We further extend the fusion-gate backbone into a timestep-conditioned diffusion denoiser with flow-matching (ODE-based) sampling and measurement-consistency guidance; while this trails the deterministic baseline in raw reconstruction quality, its multi-sample uncertainty estimates correlate strongly with reconstruction error, showing predictive value for uncertainty-aware reconstruction.

Exploring new nonreciprocal photonic platforms for controlling thermal radiation

Yiting Zhao

Mentors: Harry A. Atwater, Jr., and Simo Pajovic

Thermal radiation originates from the random motion of matter and thus pervades the world around us. Given its ubiquitousness, the control of thermal radiation has many important applications, such as radiative cooling and solar energy harvesting. Thermal emitters are constrained by Kirchhoff's law of thermal radiation, stating that the emissivity and absorptivity of a surface are equal for a given wavelength, direction, and polarization of light. Current approaches to lifting this constraint and achieving "nonreciprocal thermal radiation" emphasize materials such as InAs, which break the emissivity-absorptivity equality under a magnetic field. However, these approaches are limited, requiring a large magnetic field or only functioning for specific polarizations. Here, we explore two new platforms for nonreciprocity. The first achieves dual-polarized nonreciprocal thermal radiation using an array of silicon micropores on bulk InAs. We successfully fabricated this structure and experimentally demonstrated nonreciprocal absorptivity. The second is a theoretical thin film InAs structure which can be fabricated via spalling, a technique recently applied to III-V materials. We predict that this structure's nonreciprocal optical response differs from bulk substrates, suggesting a novel, experimentally accessible InAs platform. Our findings identify new ways to tailor thermal radiation, with potential applications in thermal management and energy conversion.

Catalytic asymmetric semipinacol-type rearrangement for the construction of all-carbon quaternary stereocenters

Junrui Zheng

Mentors: Brian M. Stoltz and Weicheng Xu

All-carbon quaternary stereocenters are common in structurally complex molecules, yet their catalytic enantioselective construction remains challenging. This project aims to develop an asymmetric semipinacol-type rearrangement in which a cyclobutanol-containing α,β -unsaturated 1,4-dicarbonyl undergoes strain-releasing 1,2-carbon migration to form a cyclopentanone bearing an all-carbon quaternary stereocenter. A modular furan-based route will provide electronically and sterically varied substrates, while reaction development will examine how chiral Brønsted-acid ion can organize the migration of two enantiotopic cyclobutane bonds. By separating productive ring expansion from competing fragmentation and establishing relationships among catalyst structure, reaction rate, and enantioselectivity, this work seeks to define a general strategy for stereocontrolled skeletal rearrangement.

Expanding past Miller-Schupp: Generating new presentation families for learning Andrews-Curtis trivialization paths using knot theory

Aaron Zhou

Mentors: Anima Anandkumar and Robert Joseph George

The Andrews-Curtis (AC) conjecture claims that all balanced presentations of the trivial group can be transformed into the trivial presentation via a sequence of moves known as AC moves. Finding trivialization paths in the AC conjecture is a difficult task due to the fact that trivialization paths can be thousands of moves long. One difficulty that arises when training machine learning models to predict the moves in an AC trivialization path is the lack of diversity in the families of presentations from which one can source data. Currently, many approaches use the Miller-Schupp family of presentations as their dataset to train their models, risking a lack of generalizability of models trained on this specific family to presentations of the trivial group at large. In order to expand the diversity of this data, we explore a connection between unknot diagrams and AC-trivial presentations by developing a pipeline to generate all unknot diagrams with at most 7 crossings and constructing their Wirtinger presentations. We then apply reduction moves to these presentations to generate a family of 4,046,634 2-generator AC-trivial presentations. We will also report on the ability of machine learning models trained on this dataset to learn the reverse mapping from presentations to knots.

Rip-stop metamaterials inspired by double gyroids

Jacqueline R. Zhou

Mentors: Angkur Shaikeea and Rosan Naik

Gyroids are ubiquitously present in nature and have inspired engineers to design efficient structures. There is a growing interest in studying gyroids to design tough architected solids. While single gyroids generally do not outperform the toughness of most other topologies, there is an opportunity to use their chiral structure in superimposing two oppositely chiral gyroids and enhance their toughness. Engineers have designed a class of textiles called rip-stop nylon that inhibits the growth of a crack, similar to how reinforcing fibers in composites can arrest cracks. However, this requires combining a second-phase material into a structure, such as a ductile fiber in relatively brittle textiles. Here, we show that double gyroids, created by translating and superimposing a second oppositely chiral gyroid onto a gyroid lattice, can be designed to inherit such properties purely from their topology, even without multiple materials. We conduct a comprehensive study by mechanically testing SLA resin 3D printed specimens to compare the strength and work of fracture of single gyroids and the family of double gyroids.

Trajectory-conditioned reinforcement learning for multimodal control of M4 robots

Warren C. Zhou

Mentors: Mory Gharib and Matthew Anderson

The Multi-Modal Mobility Morphobot (M4) is a family of reconfigurable robots whose appendages can be repurposed for aerial and terrestrial locomotion. Controlling these robots is challenging because their dynamics, actuator geometry, contact conditions, and number of controllable degrees of freedom vary across configurations and operating modes. Previous reinforcement learning work demonstrated morpholanding for fixed point-to-point objectives, but did not provide explicit trajectory tracking or unified control across ground and aerial operation. This project develops a trajectory-conditioned reinforcement learning framework for two contrasting M4 variants: ATMO, the lowest degree-of-freedom M4, and Stork, the highest degree-of-freedom M4. For ATMO, a single policy coordinates four operating modes: driving, takeoff, flight, and landing. The policy commands individual rotor thrusts, morphological reconfiguration, and wheel motion while tracking time-varying position, velocity, and heading references. The same training framework is used to train Stork policies despite its substantially larger observation and action spaces. Both the ATMO and Stork policies perform well in Isaac Lab across their targeted behaviors, demonstrating that the training approach can accommodate M4 platforms with widely different levels of mechanical complexity. The ATMO policy has also demonstrated functionality in a PX4-Gazebo simulation with actuator dynamics and interfaces that differ from the training environment. An ATMO-specific hardware deployment pipeline has been developed using PX4, a Jetson companion computer, and OptiTrack localization. Initial physical testing will focus on ATMO, after which the Stork policy may also be deployed, while the shared reinforcement learning framework can be used to train policies for additional M4 variants.

RAFT-XDIC: Deep learning-based framework for DIC measurement of mode II shear fractures

Claire Zhu

Mentors: Jin Yang, Guruswami Ravichandran, and Zixiang Tong

Digital image correlation (DIC) is a widely used optical technique in experimental mechanics for measuring full-field displacement and strain fields. Recent research has introduced deep learning-based DIC methods hoping to reduce manual user finetuning of parameters, computational effort, and obtain dense pixel-wise displacement field measurements. However, while most work has been conducted on smooth, continuous deformations, very little has been explored if these frameworks are able to resolve discontinuities and fractures.

This work presents the findings of RAFT-XDIC, a task-optimized deep learning pipeline employing the Recurrent All-Pairs Field Transforms (RAFT) architecture to yield accurate full-field DIC measurements of mode II shear deformations. A model was trained on synthetically generated speckle datasets containing piecewise rigid-body shearing mimicking a realistic mode II fracture. Comprehensive benchmarking and analysis were conducted on synthetic and experimental datasets of real speckle patterns and a dynamic earthquake rupture to evaluate model performance. The models were found to achieve sub-pixel accuracy and order of magnitude improvements over other learning-based DIC methods as well as conventional subset-based DIC.

Design and validation of a localized thermal annealing system for ultra-high vacuum copper gasket sealing

Jacob L. Zimmerman

Mentors: Richard McCarthy, Jon Feicht, Jordan Vanosky, and Melina Fuentes-Garcia

This project focuses on designing and validating a localized thermal annealing system to resolve ultra-high vacuum (UHV) sealing failures in custom copper gaskets at LIGO observatories. Initial evaluations of raw copper gaskets fabricated using basic hand tools without thermal treatment demonstrated a high failure rate, with half of the components failing to maintain a dependable vacuum seal. To eliminate this issue, we engineered a split-block heating assembly using aluminum bronze to uniformly heat the gaskets to a target temperature of 450°C. We manufactured a physical full-scale prototype using 3D printing to validate mechanical tolerances, and component fitment before physical

fabrication. Detailed mechanical drawings are being generated for manufacturing and configuration control. Our finalized system will integrate automated temperature control and ceramic base isolation to protect laboratory environments. Implementing this localized thermal treatment framework will soften the copper grain structure, allowing the gaskets to deform correctly under compression. This ensures reliable, leak-free sealing configurations required for high-precision clean and bake processes. These gaskets are one-time use and are needed to qualify parts that are to be used inside of LIGO UHV systems. Qualifying a part means that it has been cleaned to a predetermined acceptable level of off gassing.

Learning anisotropic prediction targets for flow models

Rayhan A. Zirvi

Mentors: Leonidas Guibas and Georgia Gkioxari

Flow-based generative models are trained on a hand-chosen prediction target, such as the velocity, the noise, or the clean data. This choice strongly affects training and sample quality. We introduce Anisotropic Flow Models (AnisoFlow), which replaces a fixed prediction target with a fully learnable one. AnisoFlow learns to scale the noise component of the target anisotropically, assigning a separate nonnegative weight to each direction in a jointly learned basis. A closed-form map recovers the velocity from the network's prediction, leaving the architecture, flow-matching loss, and sampler unchanged. Since recovery is exact, training optimizes the target purely to make it easier for the network to predict. We theoretically analyze the resulting prediction error and predict that the preferred noise weight in each direction scales near-linearly with data variance, in agreement with the weights learned empirically in pixel space. AnisoFlow improves FID from 1.90 to 1.55 over the JiT baseline on ImageNet 256x256, by learning only the prediction target, without representation alignment or architectural changes. In pixel-space text-to-image generation, AnisoFlow improves image quality and, through better-conditioned velocity recovery, largely removes patch-seam artifacts under reduced-precision inference. Code and models will be released.

Aluminum among the meteors: Quantifying and tracking the rising impact of space debris reentries

Alice Zou

Mentors: Randy Pollock and Paul D. Asimow

The increasing frequency of spacecraft reentries may introduce substantial quantities of Aluminum and other anthropogenic materials into the upper atmosphere, yet these events remain difficult to observe systematically from the ground. This study develops a multispectral method for identifying candidate Aluminum-rich space debris reentry events using synchronized Ultra-Violet (UV), Blue, and broadband Visible Light imaging. Aluminum has a strong, discrete atomic emission near-UV line at 396 nanometers when thermally excited; therefore, comparisons among the three channels may help distinguish the narrow-band signature of candidate artificial objects from the continuous blackbody radiation of meteors and other optical phenomena. The instrument combines a UV-sensitive industrial camera with Blue and broadband Visible Light reference cameras and adapts the Global Meteor Network detection and astrometric-calibration framework for multispectral observations. Broadband Visible Light measurements provide star-field calibration, while synchronized Blue observations serve as a spectral control. Observational validation is still required to establish the system's detection sensitivity and its ability to distinguish reentries from natural events. The long-term goal is a low-cost citizen-science network for monitoring the frequency, trajectories, and atmospheric effects of space debris reentries.

STRATA: An agentic training architecture for scientific task automation

Nicholas D. Abram

Mentor: Yisong Yue

AI agents offer a promising approach to scaling scientific data analysis by automating tasks that traditionally require substantial expert judgment. However, off-the-shelf agentic systems lack the domain expertise needed to perform the highly technical operations required by bespoke scientific data processing tasks. The performance of such agent systems can be undermined by high variability and irreproducible context. We introduce STRATA: Scientific Task, Rubric, and Agent Training Architecture, a framework for developing and evaluating reproducible scientific automation workflows via optimizing a rubric proxy that distills from expert examples and implements expert judgment in the form of executable and audible programs.

The architecture separates three roles. A Curator translates scientist-provided data, examples, and feedback into an adaptive evaluation rubric. An Implementor performs the scientific task by creating executable Python code, while an Evaluator assesses the Implementor's output against the rubric and produces grounded, actionable feedback. The Implementor then uses this feedback to generate a revised solution. The Curator receives the initial output, evaluator feedback, and revised output to determine whether the feedback produced meaningful improvement and to update the rubric for the next iteration. A private verification mechanism independently measures performance without exposing held-out evaluation information to the agents.

The architecture is designed to generalize across scientific domains and task types. By separating task execution, evaluation, and rubric development, the framework makes it possible to measure whether agent feedback produces meaningful improvement and whether the resulting guidance transfers to new tasks. The architecture therefore provides a general approach for building scientific agents that learn from expert judgment while preserving reproducibility and independent evaluation.

Phylogenetic comparative methods for the 8-Cube mouse founder transcriptome atlas

Akshay Anand

Mentors: Lior S. Pachter and Nikhila Swarna

Phylogenetic comparative methods, including Felsenstein's independent contrasts, require that samples be related by a hierarchical tree. We show that the five laboratory strains among the eight Collaborative Cross founders have extensively admixed genomes incompatible with tree-like descent and inference. The three wild-derived founder strains, representing the domesticus, musculus, and castaneus subspecies of *Mus musculus*, together with the reference strain C57BL/6J form a resolved four-taxon phylogeny suitable for expression-based tree inference. Using single-nucleus RNA-seq data from the 8-cube founder dataset, we find that the correct tree topology is recovered from gene expression across all cell types and gene sets. However, not all genes carry equal phylogenetic signal; entropy-based specificity metrics from ember effectively isolate genes with the strongest. Without such selection, we find that male and female transcriptomes diverge between strains at comparable rates while a substantial fraction of divergence is cell-type-contextual, varying by cell type rather than reflecting a uniform transcriptome-wide shift. For instance, we see divergence in the kidney can be isolated to proximal tubule and thin ascending limb epithelial cells. We extend this decomposition across all eight tissues in the 8-cube dataset to ask how cell-type divergence profiles compare across organs.

Computational study of bimetallic MOF electrocatalysts for selective nitrogen reduction reaction

Eddie I. Anderson

Mentors: William A. Goddard III and Sejun Kim

Sustainable substitutes for the Haber-Bosch ammonia production pathway have been realized through the electrochemical nitrogen reduction reaction (eNRR). Practical selectivity of the eNRR is limited by N-N bond formation and the competing hydrogen evolution reaction (HER). This project uses computational density functional theory methods to analyze how bimetallic ordering within MOF-74 influences nitrogen adsorption and to assess activation and selectivity components. Initial screening analyzed a broad range of Fe-Mo/Fe-Ni structures and narrowed our focus to mixed and separated Fe-Mo arrangements. Adsorbate-framework arrangements and electronic energies were obtained through spin-polarized PBE-D3 calculations with a 500-eV plane-wave cutoff. Sites were named after their atomic order in VESTA structure files. Mixed Fe₆ and Mo₃ sites both retain end-on N₂, while mixed Fe-Mo bridge and separated Fe sites relax to gas-like, desorbed states. On mixed structures, Fe and Mo show promise for selectivity and activation, respectively. Fe₆ is electronically favorable for N₂ adsorption by -0.40 eV, while H adsorption is unfavorable by +0.71 eV, making the site a leader in selectivity. Although Mo₃ binds N₂ more strongly (-1.36 eV) and shows that the N₂ bond length increases after adsorption, a sign of activation, this site favors H adsorption, implying HER competition. Current results show electronic energy trends. Further work is needed to determine whether Fe selectivity or Mo activation governs eNRR performance.

MACE-H-LR: Incorporating long-range interactions into equivariant Hamiltonian prediction

Jeremiah Bailey

Mentors: Marco Bernardi and Yao Luo

Accurate prediction of Hamiltonians is important for efficiently modeling the electronic properties of materials. Equivariant machine-learning models such as MACE-H can learn Hamiltonians from local atomic environments, but finite-cutoff architectures may have difficulty representing long-range interactions that extend beyond the local neighborhood. This work introduces MACE-H-LR, a physics-informed extension of MACE-H designed to predict the full Hamiltonian while explicitly accounting for both short-range and long-range contributions.

The short-range contribution is learned using an equivariant message-passing neural network, while the long-range component is incorporated through a separate physically motivated formulation based on the atomic structure and periodic simulation cell. The model predicts Hamiltonian matrix elements between atomic orbitals while preserving the rotational and geometric symmetries of the system. Developing the framework required consistent treatment of atomic-orbital indices, lattice translations, tensor dimensions, coordinate conventions, and periodic boundary conditions.

Following training, MACE-H-LR was evaluated on structures excluded from the training dataset and compared with the original short-range MACE-H model. The evaluation examined prediction errors, convergence behavior, and the relative influence of the short-range and long-range contributions. The results demonstrate the feasibility of combining equivariant machine learning with an explicit long-range treatment for full Hamiltonian prediction. This framework provides a foundation for more accurate and physically consistent electronic-structure models, particularly for materials in which nonlocal interactions significantly influence the Hamiltonian.

Towards cryo-EM structural determination of FAM136A using a nanobody-based fiducial strategy

Sandra E. Baltazar

Mentors: Rebecca M. Voorhees and Jichen Zhang

The mitochondrial membrane consists of two types of membranes: the outer membrane (OM) and the inner membrane (IM). Between these membranes exists an inner membrane space (IMS) in which various proteins, such as FAM136A, conduct their functions. Recently, it has been found that FAM136A is an IMS chaperone that helps VDACs insertion into the OM. Without this interaction, the mitochondria would not be able to maintain proteostasis, leading to stress signal responses and

human disease. This project aims to better understand this mechanism through the lens of structural biology. Cryo-EM is a tool that typically allows determination of protein structures that are the size of 100kDa or greater. Still, protein structures < 30 kDa, like FAM136A, are usually hard to analyze by cryo-EM. A developed solution is to bind a universal nanobody, NabFab, along with a binder that is specific to the target protein in order to increase its size without interfering with its functionality. We have been able to use a computational method to predict binders that are likely to bind to FAM136A. If successful, our efforts will display the accuracy of generated binders and their capabilities in solving the structure of FAM136A.

Mechanistic investigation of lanthanide reactivity and imine substrate scope in a photodriven Sm-catalyzed asymmetric ketyl-olefin coupling

Elizabeth Y. Bank

Mentors: Sarah E. Reisman, Nathan C. Friede, and Will Chen

Samarium(II) iodide (SmI₂) is a versatile single-electron reductant widely employed in organic synthesis. Although SmI₂-mediated transformations have traditionally required (super)stoichiometric quantities of reagent, recent advances have enabled catalytic Sm-mediated reactivity. Specifically, development of a photodriven Sm-catalyzed asymmetric ketyl-olefin coupling by the Reisman group has established a new strategy for catalytic samarium chemistry, providing a platform for investigating the fundamental factors governing catalytic reactivity and selectivity. This work combines mechanistic investigations with electrochemical studies, including cyclic voltammetry, to examine two complementary objectives within the Sm-PyBOX catalytic cycle: elucidating the influence of lanthanide identity on catalyst reactivity and extending the methodology to enantioselective imine-olefin cross-coupling. Screening reactivity of lanthanide-PyBOX complexes investigates how ionic radius, Lewis acidity, 4f electronic structure, and redox properties influence catalytic activity and stability. In parallel, mechanistic studies of imine substrates seek to identify the factors governing substrate activation, catalyst coordination, and suppression of competing background reactivity. Collectively, this research aims to advance the mechanistic understanding of catalytic samarium chemistry and guide the development of new enantioselective lanthanide-catalyzed reductive cross-coupling reactions.

Informing photosynthetic model parameterization using solar-induced chlorophyll fluorescence data

Nicolas W. Bartolomeo

Mentors: Tapio Schneider, Katherine Deck, and Oliver Dunbar

In global climate models, photosynthesis is an environmental process integral to accurately modeling the water, energy, and carbon cycles on a global scale. Within ClimaLand, the Climate Modeling Alliance's land surface model, the fitting of the free parameters of the modeled photosynthesis process is based on matching model-predicted values to a statistically reconstructed global dataset of a measured by-product of photosynthesis that has high bias. This project investigates an alternative dataset, specifically Solar-Induced Chlorophyll Fluorescence (SIF) satellite data, for calibrating photosynthetic parameters within ClimaLand. SIF data are collected globally by measuring photon flux during photosynthesis as plants fluoresce, making observations more direct. With Ensemble Kalman Inversion (EKI), free photosynthetic parameters are estimated from priors informed by the nonlinear model ClimaLand and TROPOMI SIF Observations, aiming to minimize the misfit function between observed and simulated values. Using daily TROPOMI SIF satellite data collected between 2018 and 2021, we calibrated photosynthetic free parameters with EKI. Using these newly found parameter values, we expect to see improvement within global photosynthesis calculations. The importance of this result is that the various environmental cycles mentioned earlier are now better informed when modeled.

Using paleoatmospheric CO₂ to constrain the age and paleoclimate of late Pleistocene Amazonian fossil flora

Nivale E. Baxendell

Mentors: Julia Tejada and Elena Stiles

Although the Amazon is a global carbon sink and regional climate regulator, little is known about how its vegetation responded to glacial–interglacial climate cycles. This knowledge is essential for predicting future environmental change in the region, yet paleobotanical records are scarce. Here, we investigate a late Pleistocene plant fossil deposit in the Peruvian Amazon. Preliminary age estimates place these fossils between 100–150 ka, an interval spanning glacial and interglacial conditions. Because atmospheric CO₂ concentrations vary with climate, reconstructed paleoatmospheric CO₂ constrains the age and climatic setting of fossil assemblages. We analyzed approximately 200 fossil leaf cuticles by combining bulk organic carbon isotope analysis with stomatal anatomical measurements from epifluorescence microscopy in a mechanistic leaf gas-exchange model to reconstruct paleoatmospheric CO₂. Preliminary reconstructions yield an average value of 480 ppm, substantially higher than the 180–280 ppm expected for the late Pleistocene. Ongoing taxonomic identification of fossil morphotypes will allow improved physiological model parameters to be assigned, producing more robust paleoatmospheric CO₂ reconstructions to better determine whether this flora represents glacial or interglacial conditions and improve our understanding of Amazonian vegetation responses to past climate change.

Deconvolving vital effects from paleoceanographic signals using stable isotopes, clumped isotopes, and metal/Ca ratios in deep-sea scleractinian corals

Jesse M. Benitez

Mentor: Jess F. Adkins

Deep-sea corals are a useful archive of thermocline and deep ocean conditions, with the potential of providing century-long records of deep ocean change at sub-decadal resolution. However, offsets from isotopic and trace-element equilibrium in coral skeletons, termed “vital effects”, have complicated the application of these archives, as well as those of other biomineralizing organisms, to paleoclimate reconstructions. Deep-sea corals are ideal candidates for the study of biological overprinting mechanisms due to the large tracer gradients within individual corals, their independence from photosynthesis, and relatively stable growth conditions. Insights into the biomineralization mechanisms of these corals might be used to develop a generalized understanding of biomineralization in other Scleractinia and marine calcifiers. In this study, we present stable isotope, clumped isotope, and trace-metal (Me/Ca) data in a suite of *Desmophyllum dianthus*, *Caryophyllia*, *Lophelia*, and *Enapsollamia* coral specimens, performing multi-proxy analyses on the same samples when possible. Both genera were selected due to their large tracer variability, wide environmental range, and availability in the fossil record. Data to come. Additionally, based on initial indications of consistent tracer variability and correlation and the suggestions of control by a single universal biomineralization process, we aim to build a biomineralization model that incorporates all three tracer classes. To do this, we incorporate Me/Ca differential equations (Chen, 2022) into a stable and clumped isotope model for biomineralization (Watkins, 2025). The outputs of this model are compared to the stable isotope, clumped isotope, and trace-metal data from the coral samples. Data to come. These findings illuminate the mechanisms of biological overprinting of isotopic and trace element composition in deep-sea coral skeletons, which is essential to the improved use of these specimens as paleoceanographic archives of temperature, pH, carbonate chemistry, and other environmental parameters.

Performance evaluation of a kinetic inductance traveling wave parametric amplifier

Clayton G. Bishop

Mentor: Sunil Golwala

For the continued search for dark matter (DM), low-temperature phonon detectors are a promising technology for probing the sub-GeV DM parameter space. Kinetic Inductance Detectors (KID) are sensitive to energy depositions in a substrate through athermal phonons generated by a potential interaction. Kinetic Inductance Traveling Wave Parametric Amplifiers (KI-TWPA) improve energy resolution through their high gain and nearly quantum-limited noise, thereby maximizing the signal-

to-noise ratio. In this work, we outfitted a dilution refrigerator with proper components to operate a KID and KI-TWPA. After cooling down, we measured the gain and bandwidth of the parametric amplifier and determined its added noise using the Y-factor method.

Controlling mist formation in droplet impact using megasupramolecules

Paxton N. Boyce

Mentors: Julia A. Kornfield and Satvik Verma

Controlling mist formation in liquid sprays is essential in many applications, including fuel safety. In the event of a crash, flammable hydrocarbons are sprayed outside the confines of the engine, generating an explosive air-fuel mixture. It has been previously demonstrated that adding dilute concentrations of ultralong polymers to fuel sprays decreases flammability by decreasing mist formation; however, the mechanisms by which these polymers decrease mist formation are not fully understood. To investigate these mechanisms, we study the effects of adding megasupramolecules to polyalphaolefin lubricant in the droplet impact flow configuration. Using high speed images of this impact, we quantify secondary droplet size distributions, droplet and bead fates, and local kinematics of pinchoff events. Secondary droplet sizes follow a gamma distribution up to 0.01 wt% megasupramolecules until transitioning to a bimodal distribution. This transition is due to the emergence of the beads-on-a-string morphology and subsides at higher concentrations due to longer filament lifetimes. Bead fates confirm this, showing that beads deposit on the solid substrate at higher megasupramolecule concentrations before they can pinch-off to form secondary droplets. All analyzed ligaments lie in the Rayleigh regime of the Ohnesorge diagram, allowing us to study their breakup behavior using simpler flow configurations, such as Rayleigh jets. This analysis confirms that megasupramolecules are effective in decreasing fine mist formation by increasing deposition of beads. Future studies should investigate how ultralong polymers change splashing behaviors, as droplet impact experiments show a transition from corona splash to levitating lamellar breakup and total deposition at higher concentrations.

Synthesis and characterization of transition metal complexes for molecular qubit applications

Carolyn S. Brown

Mentors: Theodor Agapie and Fernando Guerrero

The field of quantum information science leverages quantum mechanical phenomena, such as superposition and entanglement, to advance computing and sensing. The basic unit of quantum information is the quantum bit, or "qubit," which can exist in a superposition of both the 0 and 1 state. Quantum bits exploit superposition to enable sensing with precision beyond classical instruments, and molecular systems offer a synthetically tunable platform towards this goal. While open-shell copper(II) and cobalt(II) complexes have demonstrated anion sensing measurable via electron paramagnetic resonance (EPR) spectroscopy, vanadium(IV) qubits remain comparatively unexplored despite favorable spin coherence properties. This work reports the synthesis and characterization of a new vanadyl qubit candidate supported by a *bis*-triazinyl bipyridine (**BTBP**) ligand, whose open equatorial coordination site allows variable ligand binding. Complexes of the general form VO(BTBP)(OTf)₂(L) were synthesized and characterized by ¹H NMR and continuous-wave (CW) EPR spectroscopy. CW-EPR revealed measurable shifts in electronic structure upon equatorial ligand exchange, indicating that this platform is sensitive to bound analytes. Future work will apply pulse-EPR to assess spin relaxation dynamics and explore pH-sensing applications of this system.

Investigating the molecular and neural mechanisms of social behavior in *Drosophila*

Kamarya A. Buckley

Mentors: David J. Anderson, Yousef Ouadah, and Ziqing Zhong

Social behaviors such as aggression and courtship are regulated by neural circuits and pheromone signaling, yet the molecular mechanisms underlying their evolution remain poorly understood. *Drosophila santomea* provides a valuable model because males exhibit increased same sex courtship and produce reduced levels of cis-vaccenyl (cVA), a pheromone that promotes aggression and suppresses male-male courtship in most *Drosophila* species. This study investigates the cellular basis of cVA production by characterizing the ejaculatory bulb, the organ responsible for cVA biosynthesis,

using single-cell RNA sequencing (scRNA-seq). Current work focuses on optimizing tissue dissociation and validating cell viability to generate high-quality single-cell suspensions for sequencing. These data will ultimately identify candidate genes involved in cVA biosynthesis, with particular interest in epithelial cell populations.

In parallel, in *Drosophila melanogaster*, we examined the neural circuitry underlying aggression through an epistasis-based optogenetic approach by activating common aggression-promoting (CAP) neurons while silencing the downstream neuron SMP054; we tested whether SMP054 is required for aggression-related behaviors. Preliminary behavioral observations suggest SMP054 may not be required for CAP-mediated aggression and approach behaviors. Together, these complementary approaches will advance our understanding of the molecular and neural mechanisms that shape the evolution of social behavior in *Drosophila*

PySimSat: An open-source Python package for simulating satellite X-ray spectra

Michael Burch

Mentors: Fiona A. Harrison and Gaurav Waratkar

We present PySimSat, an open-source Python package designed to simulate X-ray spectra and evaluate instrument performance for concept satellite missions. Existing mission performance simulations combine mission specific software with ray tracing toolkits such as GEANT4 to model detector response and observed spectra. However, these simulations are often computationally expensive, complex, and time-intensive, with some simulations taking weeks or even months to complete. Additionally, these simulations are generally developed for a specific mission and are not easily adapted to new mission concepts. We developed PySimSat to provide fast, simple performance simulations of concept missions. Given the characteristics of the concept mission, including satellite geometry, detector properties, optics, and orbital configuration, PySimSat generates satellite response files and simulates the observed background and source spectra. We use the Burst Alert Telescope (BAT) and AstroSat as case studies to validate PySimSat by comparing PySimSat simulated spectra with actual observations. We then apply PySimSat to optimize a CubeSat mission designed to constrain the corona temperature of active galactic nuclei (AGN) and investigate the processes responsible for X-ray emission in the corona.

Grain boundary evolution for a 3D X-ray diffraction inversion algorithm

Eitan E. Camacho

Mentors: Kaushik Bhattacharya and Carter Cocke

We extend and compare two level set methods for evolving the boundaries of many-phased systems: a multiple level set approach and the Voronoi implicit interface method. The multiple level set approach assigns a level set to each phase and advects them independently. A corrector step is required to fill voids and resolve overlapping regions. The Voronoi implicit interface method tracks any number of phase boundaries with a single unsigned distance function and an indicator function. While the initial Voronoi implicit interface method was implemented on a Cartesian mesh, we extend the algorithm to unstructured tetrahedral meshes for advection due to an externally given, time-independent velocity field. Unstructured meshes provide additional flexibility crucial to working with complex geometries. In both cases, the level sets are reinitialized using the fast marching method, an efficient algorithm for finding distances. The methods are used as part of a forward solve in an inversion algorithm that recovers intragranular strain and orientation fields from high-energy X-ray diffraction data. While the approaches appear vastly different, we note that they have the same time complexity and early stage computational results support their near equivalence.

Flow Monte Carlo tree search

Jason H. Childs

Mentors: Soon-Jo Chung and Min Kim

This project aims to implement Conditional Flow Matching (CFM) on Monte Carlo Tree Search (MCTS) distributions. MCTS algorithms utilize loops of selection, expansion, simulation, and backpropagation to determine a most optimal set of behaviors to reach a desired outcome. CFM models the dynamics evolving one distribution to the next distribution to normalize data. Its use in this project is to evolve

a root state distribution to its successor state distribution. Flow Monte Carlo Tree Search (FMCTS) is a hybrid algorithm for more efficient planning where the entirety of the MCTS tree no longer needs to be built. Instead, the tree search is compressed into a neural network. The benefits of this new system design aim to improve speed, improve batch inferences, and create smoother actions within the running system. Preliminary data will be discussed in presentation.

Investigating the origin of seismic clustering in the Salton Trough

Claire D. Fagan

Mentors: Joann M. Stock and Allen Husker

California's Salton Trough (ST) region experiences oblique extension between the Pacific and North American Plates, geothermal and magmatic activity, changes in hydrological loading, and high seismicity. We aimed to identify causes of seismicity in this geologically complex region. We first used Southern California Earthquake Data Center (SCEDC) catalog data from 2023-2026 to identify seismicity and separate events into (1) swarms, (2) mainshock-aftershock sequences, (3) mixtures containing attributes of both (1) and (2), and (4) background seismicity. We correlated borehole well data from ST geothermal fields with the seismic data, finding that swarms and mixtures occur within a 3 km radius of geothermal injection wells twice as often as mainshock-aftershock sequences and background seismicity. Swarms and mixtures within the injection radius occurred at shallower depths and lower magnitudes. Many clusters outside the injection radius were planar. To investigate these clusters, we mapped new 3D faults in the ST using relocated seismicity from the SCEDC catalog and compared the clusters to all updated regional faults. Overall, we interpret that seismic clusters within the injection radius of geothermal wells and fields occur due to increased pore fluid pressure generated by geothermal activity, while clusters outside of injection radius occur via strike-slip fault friction.

Characterization of germinal center B cell-derived antibodies

Jordan L. Foster

Mentors: Pamela J. Bjorkman, Alexander Cohen, and Jennifer R. Keefe

An antibody response is a key component of the immune system in which B lymphocytes produce antigen-specific antibodies that bind to and help eliminate pathogens. Prior work in the Bjorkman lab has shown evidence indicating that Germinal Center (GC) B cells arising from secondary GC immune responses (from boost/second dose) have a greater capacity to recognize a broader range of antigens, including antigens that were not directly encountered during the initial immunization. In contrast to GC B cells generated during the primary immune responses (from prime/ 1st dose) appear to be more restricted in their specificity and recognize the cognate antigen or the antigen used for initial immunization. My project sought to understand the restricted specificity of the GC B cells. It involved the cloning and expression of monoclonal antibodies (mAbs) from GC B cells isolated after prime and boost immunizations with model viral antigens, specifically sarbecovirus receptor binding domain (RBD) nanoparticles and influenza hemagglutinin (HA) nanoparticles to evaluate antigen specificity of the mAbs through the testing of their binding to cognate and non-cognate antigens, as well as their Ab function using pseudovirus neutralization. Preliminary results will be discussed.

Decoding behavioral states and transition dynamics in human driving task

Deborah R. Fu

Mentors: Markus Meister and Yingxi Jin

The naturalistic approach to human neuroscience study involves the study of human neural and behavioral data during complex tasks. A complex task is defined as one that requires the sequencing of functioning distinct episodes, or "microtasks". Here, we report the behavioral state boundaries of human subjects engaged in a naturalistic, goal-directed complex task: driving a simulated car around a custom figure-8 race track. SVM analysis reveals that microtask decisions become decodable before the microtask boundary is reached. Behavior along the track is modeled as a recurrent switching linear dynamical system. Shared discrete dynamical states and state transitions along the track are identified for neural correlation analysis.

Employing novel spectroscopic techniques in tracking biomarkers through coalification

Gabriella G. Giampietro

Mentors: John M. Eiler and Annie Edwards

Coal is a heterogeneous material made up of carbonaceous components derived from ancient flora. Quantifying biomolecule abundance in coals opens new avenues into understanding terrestrial paleoclimate at the time of deposition. However, current methods of quantifying these biomolecules are largely destructive, posing problems when analyzing sample-limited materials. Spectroscopy offers a non-destructive alternative for analyzing the geochemistry of fossil leaves and coals. We present new insights in the application of a range of traditional and novel spectroscopic methods, including bond selective fluorescence-detected infrared-excited (BonFIRE) spectroscopy, to the study of coal geochemistry. ATR-FTIR spectroscopy and Raman spectroscopy were conducted on coal rank samples and chemical standards including chlorophyll, lignin, and cellulose. Resulting spectra were analyzed and compared to chemical spectra of known compounds in order to predict visible signals when using novel spectroscopy techniques. Following literature review, BonFIRE spectroscopy was selected due to its potential to examine chemical specificity with greater sensitivity than traditional vibrational spectroscopy, making it well-suited to fossil applications. BonFIRE was then performed on the sample set and analyzed compared with established spectroscopic libraries. Preliminary results identify challenges of using BonFIRE in coal geochemistry, including selecting appropriate dyes to stain target molecules, navigating around machine blind spots, and sample preparation methodology. There is potential in applying BonFIRE to distinguish between molecules contributing to broad composite bands, such as aromatic carbon-carbon skeletal vibrations at 1593cm⁻¹. Future directions include conducting additional novel spectroscopy methods on the sample set.

Why do LLMs misdescribe themselves?

Zakiriya H. Gladney

Mentor: R. Michael Alvarez

Language models describe their own dispositions consistently, yet those self-reports predict their behavior poorly. We ask where this say-do gap comes from. Coherence requires preserving a hierarchy—a broad trait above specific tendencies, above the behaviors that express them—and we test three mechanisms that could distort or discard it. First, a Transformer represents everything in a single flat (Euclidean) space, too cramped to keep the branches of a deep hierarchy apart; we test whether a curved (hyperbolic) space, which holds it with far less distortion, closes the gap. Second, it may have been discarded during training, since next-token prediction never demanded it. Third, self-report and behavior may be produced by separate mechanisms, read from directions that no change of representation can align. To tell these apart, we fine-tune two adapters on a single base model, differing only in curvature—Euclidean (LoRA) and hyperbolic (HypLoRA)—under an objective that rewards agreement between self-report and behavior shifts. Curving the space tests the geometric mechanism; complementary probes—linear decodability, causal steering, and a meaning-first (JEPA) representation—separate the other two when geometry alone doesn't close the gap. If curving restores coherence, the gap is geometric; otherwise, they identify the responsible cause. This localizes where—and whether—coherent self-models can be trained into current LLMs.

Synthesis of complex ring systems via a [2+2] cycloaddition followed by a base-induced rearrangement

Elise R. Hastie

Mentors: Brian M. Stoltz and Chloe Cerione

Dimeric complex ring systems are observed in sesquiterpenoid dimers that exhibit potent anticancer, anti-inflammatory, and neuroprotective properties; however, synthetic access to these molecules remains limited due to their structural complexity and lengthy synthetic routes. Building upon previous research in the Stoltz group, we test the viability of a [2+2] cycloaddition and base-induced cascade approach for rapid access to various complex ring systems. To assess the scope of this approach, we have dimerized multiple simple butenolide substrates via [2+2] cycloaddition and subjected the resulting dimers to sodium methoxide to observe the resulting rearrangements. Additionally, we are working toward the synthesis of chloranthalactone A, which we expect can be dimerized and subjected to base treatment to yield chloranthalactone F. All chemical structures were confirmed by nuclear

magnetic resonance (NMR) and/or single-crystal X-ray diffraction (SCXRD). Ongoing work will optimize yield and expand upon the substrate scope established by this work and previous studies conducted by the Stoltz group.

Building a minimal toolkit for multicellularity: Investigating T cell regulation through synthetic circuits

Kaedan A. Hendrix

Mentors: Michael B. Elowitz, Ethan Richman, and Jacob Parres-Gold

Multicellular organisms utilize intricate biomolecular circuits to coordinate complex behaviors, such as the human immune system's response to infection. Isolating and mastering the minimal design elements of these systems clarifies the functional trade-offs inherent to various circuit architectures. Revealing these design principles is vital to understanding multicellular communication networks and could enable the use of programmable cell circuits for advanced biomedical practices. A critical node and example of this cellular decision-making is the delicate balance between conventional effector T cells, which drive immune defense, and regulatory T cells (Tregs), which actively suppress expansion to protect "self" antigens and prevent autoimmunity.

To uncover the mechanics of this balance, a computational ODE model of alternative circuit designs for T cell regulation by Tregs was constructed in a Python environment. This model is able to systematically explore whether each distinct design can achieve a desired function: permitting conventional T cells to expand in the presence of exponentially increasing stimulation, while remaining strictly controlled under low antigen induction. The model evaluated a multi-axis design space, assessing T cell proliferation drives such as private signal sensing or direct antigen-presenting cell contact. Furthermore, distinct suppression mechanisms including direct private-receptor-facilitated killing, proliferative arrest, and protease-mediated signal cleavage were tested. Computational screening revealed that effective regulation is primarily driven by delayed Treg activation, with direct killing and protease-mediated reduction computationally suggested to be the most robust strategies. By constructing the candidate circuits in mammalian cell lines with synthetic protein components, we are able to test the results of this model *in vitro*. Ultimately, this approach will allow the direct characterization of communication circuits in live mammalian cells, revealing important design principles.

Characterizing Andrews-Gordon summands in terms of subsets of cylindric partitions

Michelle Holland

Mentors: Tony Yue Yu, Weihong Xu, and Yifeng Huang

The Andrews-Gordon identities are Rogers-Ramanujan type identities, meaning they are q -series identities equating infinite sums with infinite products. In particular, they are direct generalizations of the classical two Rogers-Ramanujan identities, valued for their applications across many fields of math and physics. This project aims to enrich an existing proof of the Andrews - Gordon identities by giving an interpretation of each summand individually, as the original proof does not present a direct term-by-term correspondence. In that proof, the sum side is related as a whole to a generating function of cylindric partitions via B-paths, recovering the product side from a known product formula for this generating function. Our aim is to relate each summand individually to a particular subset of cylindric partitions, making this established connection more granular and structurally transparent. Time permitting, we also aim to eliminate the use of B-paths in the description of these subsets, which would allow for multi-row generalizations, and so would potentially lead to proofs of the conjectural HJO identities.

The impact of snow accumulation on the Southern Ocean overturning circulation

Jasmine E. Hollis

Mentors: Andrew F. Thompson and Sarah Zhang

The world's oceans are connected by the global overturning circulation, which regulates the storage of heat and carbon dioxide in the deep ocean (Pellichero et al., 2018). Water mass transformation (WMT) refers to the processes by which water masses are lightened and densified (Abernathey et al., 2016); WMT is a key process in setting the strength of the Southern Ocean circulation's meridional

component. Previous studies have suggested that WMT rates in this region are dominated by sea ice growth and melt (Pellichero et al., 2018; Abernathey et al., 2016), yet little is known about the contribution of snow on sea ice to this process due to the complexity of snow-sea ice-ocean interactions and scarce observations (Webster et al., 2018). This study uses ECCO version 4, release 5 ocean state estimate solutions to quantify the contribution of snow accumulation on sea ice to WMT rates and how this contribution is affected by evolving sea ice extent (SIE). Preliminary results suggest a peak in snow contribution to WMT rates in November of about -11.5 Sverdrups in the region of transition between sea ice formation and sea ice melt, and suggest that the snow contribution may account for up to 12% of the total sea ice freshwater contribution to WMT rates. Initial comparisons of WMT rates between positive and negative SIE anomaly years do not indicate a significant difference. This apparent similarity will be investigated in future work.

Removing AGN contamination from millimeter-wave datasets to refine multiprobe triaxial analysis of galaxy clusters

Chandrasekhar Kappagantula

Mentors: Jack Sayers, Adriana Gavidia, and Harshda Saxena

The standard hierarchical model indicates that the formation of galaxy clusters, the largest gravitationally-bound structures known so far, are modelled better by a triaxial shape, rather than a simpler, spherically-symmetric model. Thus, parameters of a galaxy cluster can be biased if modelled spherically, and a triaxial fit requires deep, multiprobe analysis with various datasets. The Cluster Heritage project with XMM-Newton - Mass Assembly and Thermodynamics at Endpoint of structure formation (CHEX-MATE) sample contains 122 galaxy clusters which have been triaxially modelled, but one of the observables used to constrain the line of sight (LOS) and geometry, the Sunyaev-Zeldovich (SZ) effect, is at risk of contamination from Active Galactic Nuclei present in the foreground or at the center of these clusters, as a Brightest Cluster Galaxy (BCG). The SZ effect, a result of inverse-Compton scattering of CMB photons gaining energy from high-energy electrons in the cluster, is used to map and constrain intracluster pressure; thus, AGN contamination can bias pressure profiles and thus the triaxial geometry determined from the fitting. This project contains two key steps to mitigating this contamination: 1) Locate AGN contamination using the NRAO VLA Sky Survey (NVSS) and Very Large Array Sky Survey (VLASS) radio catalogs, extrapolating the fluxes of any located targets at 1.4 GHz and 3 GHz respectively to 90 GHz, roughly estimating the flux of the contaminant in the SZ effect map, and 2) Creating masks for the contamination on the SZ effect maps and fitting the newly masked observables, to evaluate if the bias is reduced. Despite the practical difficulties of locating the AGN contamination and evaluating which masks are best suited to optimally reduce fitting bias, the general procedure has proven to be effective for certain clusters, such as PSZ2 G073.97-27.82 (Abell 2390).

Strain-engineering of MoS₂/graphene/metal heterostructures

Jiri Kataman-Kustwan

Mentors: Nai-Chang Yeh, Jen-Te Chang, and Wen-Hao Chang

Excitons hosted in van der Waals heterostructures of graphene and the transition metal dichalcogenide MoS₂ are sensitive to substrate strain, which reshapes their band structure and thereby the exciton dynamics relevant to optoelectronic devices. This project combines atomistic modeling with a fabrication study of a MoS₂/graphene/gold-nanopillar stack. LAMMPS molecular dynamics relaxation supplies the strain field, which feeds a tight-binding model that yields the strain-modified band structure. Benchmarked against published scanning tunneling microscopy data for graphene on Au, the Abell-Tersoff potential reproduced lattice constants to within 1.4% and 0.4%, though moiré periodicity deviated by 14.9%; pillar-like nanostructures produced the highest localized strain near their edges. In parallel, plasma-enhanced chemical vapor deposition (PECVD) was surveyed on copper foil and nickel-on-silicon across CH₄/H₂ ratios and plasma powers, characterizing morphology by atomic force microscopy (AFM) and scanning electron microscopy (SEM) and carbon phase by Raman spectroscopy. Cu yielded no graphene, but amorphous carbon deposited downstream of the plasma indicated transported carbon species favor growth; guided by this, Ni/Si at low CH₄/H₂ ratios produced clear graphene signatures. Together, these results identify Abell-Tersoff as the preferred potential for graphene/metal strain modeling and downstream, carbon-lean PECVD on Ni/Si as a scalable route to strained heterostructures.

Isolating and quantifying the pre-targeting complex of the GET pathway in *Aspergillus fumigatus*

Avery C. Love

Mentors: William M. Clemons and Vida Robertson

Aspergillus fumigatus is a deadly infectious fungus that has evolved resistance to existing antifungal drugs. Retro compounds have been shown to enhance the effectiveness of fungicides on drug-resistant fungi, making them promising candidates for combating antimicrobial resistance. Retro compounds in human cell lines inhibit retrograde transport in a non-cytotoxic manner by disrupting the Guided Entry of Tail-anchored proteins (GET) pathway, which transports tail-anchored proteins to the endoplasmic reticulum. Specifically, Retro compounds lock the pre-targeting complex of human Get3/4 and Bag6 in a stable conformation that prohibits tail-anchors from being transferred to the main chaperone on the GET pathway, Get3. Retro compounds have also been demonstrated to interact with the *A. fumigatus* Get3, however, the mechanism of action of Retro compounds on the fungal GET pathway is undetermined. If the mechanism of action is conserved across eukaryotes, the inhibition of retrograde transport in filamentous fungi could provide a novel, potent target for antimicrobial therapeutic development. Thus, we purified AfGet3 and AfGet4/5 to better understand the *A. fumigatus* pre-targeting complex and aim to quantify the effect of Retro compounds on the fungal GET pathway

Improving stability of peptide displaying Tobacco Mosaic Virus-Like Particles

Dzianis Mahila

Mentors: Gözde S. Demirer and Andres Heredia

Tobacco Mosaic Virus Virus-Like Particles (TMV VLPs) are a promising evolutionarily adapted, yet non-infectious plant delivery system due to their unique physiochemical characteristics. TMV VLPs consist of a single coat protein (CP) that assembles into nanorod particles, which can be modified by fusing peptide cargo to different regions of the CP. However, this challenges the proper assembly of TMV VLP and can potentially cause their aggregation. This work aims to engineer the VLP scaffolds that best tolerate the peptide addition to their surface, by investigating: (1) how CP net charge and (2) stoichiometry of fusion CP within a VLP affects assembly and aggregation. The latter is explored using the Foot-and-Mouth Disease 2A skip-sequence, whereby both wild-type-like and fusion coat proteins can be co-expressed to assemble together into chimeric TMV VLPs. Our VLP characterization data supported previous findings that altering pI will disrupt TMV nanorod assembly, as our more basic CP fusions exhibited a hypothesized truncation event of their SpyTag cargo. By integrating these two approaches to reduce cargo occupancy and improve steric considerations, while also making CP charge closer to wild-type, this work demonstrates a promising strategy to re-design the scaffold of a natural plant pathogen for plant benefit.

Time-resolved domain dynamics of microwave-driven quantum quenches in 3D quantum magnet LiHoF₄

Vasilisa K. Malenkiy

Mentors: Thomas F. Rosenbaum and Daniel Silevitch

We demonstrate controlled movement of the ferromagnetic-paramagnetic (FM-PM) phase boundary in the model Ising magnet LiHoF₄ and probe quenched domain dynamics using a time-resolved ringdown technique. Hyperfine electronic-nuclear (¹⁶⁵Ho spin) coupling increases the critical field. Resonantly pumping the E₃₂ Ho³⁺ electronuclear transition (~3 GHz) redistributes state populations, providing intrinsically fast control. Rather than ramping a transverse magnetic field over many minutes, we move the phase boundary, driving the system out of equilibrium in microseconds. We use a fast microwave circuit and loop-gap resonator loaded with a pure LiHoF₄ crystal at 30 mK, and measure the evanescent response after switching a detuned probe field off. Down-mixing the resonator's free decay measures the resonant frequency f_0 and quality factor Q , yielding the susceptibility of the sample, $\Delta f_0/f_0 \approx -\eta\chi'$ and $\Delta Q^{-1} \approx \eta\chi''$. Mapping f_0 across pump power and transverse field traces the

phase boundary $P_c(B_z)$. Gating the drive at controlled rates quenches the system across the PM→FM boundary, and the resulting $\chi'(t)$ and $\chi''(t)$ time series - tracking collective-mode softening and domain wall dissipation - will resolve how topological domain-wall density scales with quench rate, constituting the first experimental test of Kibble-Zurek scaling in a macroscopic three-dimensional quantum magnet.

Development of calibration sources and data acquisition software for hyperfine spectroscopy

Liliana L. Martin

Mentors: Nick R. Hutzler, Madison Howard, and Roy Ready

Hyperfine spectrometers facilitate the rapid acquisition of high resolution, broadband spectroscopy of cold radicals. Expanding the spectrometer's calibrated range requires stable single frequency sources to serve as absolute wavelength references within each calibration range. We constructed a homebuilt, Littrow configured external cavity diode laser (ECDL) at 522.09 nm, whose alignment supports construction of ECDLs at other wavelengths by replacing the laser diode and aspheric lens. Decreasing systematic noise within calibrated ranges requires modification of the calibration and intensity versus wavelength mapping algorithms. Implementing these algorithms requires acquisition of raw images from the spectrometer camera using Python software that interfaces with the camera driver through wrapped C++ functions. Expanding the spectrometer's calibrated bandwidth while reducing systematic noise increases the range of molecules whose electronic band structure can be studied at sufficiently high resolution to perform future quantum metrology experiments.

Single-step megabase addition and replacement in bacteria via conjugative CRISPR-associated transposition

Ian M. Mathews

Mentors: Kaihang Wang and Raymond Zhang

Building chimeric genomes by combining diverse genetic repertoires across distinct bacterial species enables the melding of complex biological functions. However, existing custom genome construction methods depend on double-strand breaks or labor-intensive *in vitro* assembly, making them resource-intensive, time-consuming and unscalable for transferring megabase-sized natural DNA sourced from a diversity of microbial genomes.

To enable genome rewriting on the megabase-scale with natural DNA, we developed a suite of technologies that couple conjugative transfer with CRISPR-associated transposition (CAST) to deliver and integrate targeted genomic fragments in a single step. We evaluated Additive Conjugative-CAST Engineering (ACE) and its donor-engineering extension, Orthogonal ACE Sequential Integration System (OASIS) to transfer 80 Kb and 120 Kb segments of *Pseudomonas protegens* Pf-5 encoding for diverse biosynthetic gene clusters downstream of the *glmS* locus in *Pseudomonas putida* KT2440.

We further benchmarked Replacements with ACE (RePLACE) with cross-genus megabase replacements of 180 Kb, 680 Kb, and 1.08 Mb segments from an attenuated *Shigella flexneri* donor (CFS100) with the corresponding segment of a reduced-genome *Escherichia coli* recipient (MDS42). These results and characterizations establish conjugative-CAST systems as a scalable, programmable platform for rapid, multi-megabase synthetic genome expansion.

Characterizing astrocyte-like glia morphology in the *Drosophila* larval antennal lobe

Brianna N. Newton

Mentors: Carlos Lois and Aubrie L. De La Cruz

The brain consists of two primary cell types, neurons and glia, that work together to maintain normal neural activity. Similar to humans and other animal models, the *Drosophila melanogaster* nervous system is composed of neurons and glia, with the largest glial subtype being astrocytes.. Although neurons are well characterized for their structure and function, few studies have investigated the physiology of local glial cells within a subset of neurons. In particular, the larval *Drosophila* olfactory system consists of 21 olfactory receptor neurons (ORNs), each of which projects to exactly one glomerulus in the antennal lobe (AL). To investigate the morphology of astrocyte-like glia, we utilized

the TRanscellular ACTivation of Transcription (TRACT) system to visualize neuron-astrocyte contact. This genetic tool uses Gal4-UAS and lexA-lexAop systems to drive expression of a ligand and receptor in neurons and astrocytes, respectively, in third instar *Drosophila* larvae. Our next step is to induce apoptosis in astrocytes interacting with the ORNs to understand the local cellular response to astrocyte death. Following dissection, larvae brains were immunostained and scanned using confocal microscopy. While our study characterized astrocyte morphology in the larval AL, future directions include investigating astrocytes in other conserved larval brain areas or studying other local glial cells.

Molecular characterization of mRNA quality control in humans

Ifediora P. Nwakuche

Mentor: Isaac Fianu

RNA polymerase II (Pol II) transcription is coordinated with RNA processing and quality-control mechanisms that regulate transcript stability and transcription termination. Decapping protein 2 (DCP2) removes the 7-methylguanosine cap from RNA, generating a 5'-monophosphorylated RNA substrate that can be degraded by 5'→3' exoribonucleases [PMID: 21935889]. Although primarily associated with cytoplasmic mRNA decay, DCP2 has also been implicated in nuclear RNA quality control and premature Pol II transcription termination [PMID: 14749774; PMID: 22483619]. Previous cellular studies suggest that DCP2 may coordinate with the 5'→3' exoribonuclease XRN2, potentially coupling RNA decapping and degradation to transcription termination; however, the biochemical basis of this relationship remains unclear.

This study aims to define the functional relationship between DCP2 and XRN2 using purified recombinant proteins and biochemical approaches. Preliminary pull-down assays examining XRN2 and the DCP2 cofactor DCP1A showed no detectable direct interaction under the conditions tested. Recombinant human DCP2 has since been successfully expressed and purified, enabling direct investigation of DCP2-XRN2 coordination. Ongoing studies will use pull-down assays to test complex formation and decapping and RNA degradation assays to determine whether DCP2 activity facilitates XRN2-dependent RNA degradation.

Developing chemical tools for the inhibition of chondroitin sulfate glycosaminoglycan sulfotransferases to enhance neuronal regrowth

Martin B. Panicacci

Mentors: Linda C. Hsieh-Wilson, Cassandra N. Cancemi, and Sophia Gonzalez

Chondroitin sulfate proteoglycans (CSPGs), ubiquitous in the extracellular matrix of the central nervous system, consist of core proteins modified with chondroitin sulfate (CS) glycosaminoglycan (GAG) chains. CSPGs are diversified through the patterned sulfation of their GAG chains by sulfotransferases, which is known to influence interactions with cytokines, receptors, and growth factors. Notably, the CS-E sulfation motif has been associated with potent inhibition of neuronal recovery, thus the disruption of CSPG expression may yield insights towards the development of therapeutics. However, tools to interrogate the role of specific sulfation motifs are not widely accessible. Previously, oxindole and cyanoacrylamide scaffolds have been investigated as inhibitors of Chst15, the sulfotransferase responsible for the CS-E motif. Preliminary screening studies have identified Rottlerin as a potent inhibitor of Chst15. Towards the definition of a structure-activity relationship and tuning its activity, we have synthesized a fragment of Rottlerin and derivatives exhibiting substitutions known to increase potency against Chst15 in established scaffolds. In the future, these will be evaluated by a fluorescence-based enzyme-coupled assay and disaccharide analysis of treated Neu7 astrocyte cell cultures. Together, this work aims to further understand Rottlerin as a Chst15 inhibitor, expanding our library of GAG sulfotransferase inhibitors.

Disk-wind diagnostics from the 2025 outburst of black hole transient IGR J17091-3624

Sarah L. Pinto

Mentors: Fiona A. Harrison and Shina Adegoke

Black hole X-ray binaries (BHXBs) undergo outbursts when the mass accretion rate onto the black hole increases dramatically, producing a rapid increase in X-ray luminosity. BHXBs frequently exhibit signs of emitted disk-winds during the soft-state of an outburst. The BHXB candidate IGR J17091-3624

underwent a failed outburst in 2025 that showed evidence of highly ionized iron XXV or iron XXVI lines, indicative of disk-winds, despite the outburst never reaching the soft-state. This outburst was monitored by several X-ray missions, including NuSTAR and XRISM. While absorption lines were detected by NuSTAR, these lines appear to be missing from concurrent observations with XRISM, despite its higher spectral resolution. We aim to identify and characterize iron XXV or iron XXVI lines in NuSTAR and XRISM data. Then, we use Monte Carlo simulations to determine the statistical significance of these lines and compare the detections of the two telescopes. This comparison will help clarify the relative sensitivity of NuSTAR, with its larger effective area, and XRISM, with its higher spectral resolution but smaller effective area. Our results will also provide insight into the underlying mechanisms driving disk-winds in outburst.

Infrared variability studies of Active Galactic Nuclei with SPHEREx

Amelie E. Rajab

Mentors: Fiona A. Harrison and Elias Kammoun

The recent release of Spectro-Photometer for the History of the universe, Epoch of Reionization, and Ices Explorer (SPHEREx) data has ushered in an exciting time for the study of variability in Active Galactic Nuclei (AGN). These are the supermassive black holes that lie at the centers of galaxies and consume tremendous amounts of surrounding matter. SPHEREx presents a unique way of studying these objects, as it is the first all-sky survey in the near infrared (spanning $\lambda = 0.75\text{-}5\ \mu\text{m}$ and 102 wavelength channels), and it surveys the North and South Ecliptic Poles daily, amassing tens of thousands of observations of the same patch of sky. SPHEREx also makes use of a Linear Variable Filter (LVF) in order to capture the same sources in different wavelengths, providing thorough spectra. In this paper, we present a comprehensive data reduction process of SPHEREx data on a sample of AGN with the use of the SPExPI pipeline. The results consist of background-subtracted and error-handled spectra and light curves that span 102 wavelength channels and 10 months of daily observations. These products provide the foundation for future studies of near-infrared AGN variability with SPHEREx.

Using the Electrophoretic Mobility Shift Assay (EMSA) to test direct DNA binding in dArc1 autoregulation

Faith M. Rosas

Mentors: Kai G. Zinn and Peter H. Lee

Autoregulation is a regulatory mechanism in which a gene produces a protein that binds to its own regulatory DNA sequence, thereby regulating its own expression. The *Drosophila* Arc protein (dArc1) was found to bind specific regions of the genome, referred to as "peaks," which were identified through ChIP-seq analysis. These findings suggest that dArc1 may positively regulate its own transcription by activating expression of its own gene. This project investigates dArc1 autoregulation by testing whether it binds to regulatory DNA sequences that were identified using CHIP-seq analysis. Genetic studies in *Drosophila* were used to evaluate the role of dArc1 in regulating its own expression *in-vivo*. Candidate DNA sequence binding sites were cloned into reporter constructs containing dArc1 protein for S2 cell transfection assays to determine if dArc1 binds to candidate regulatory DNA sequences regulating its own gene expression. Electrophoretic mobility shift assays (EMSAs) were performed to assess direct binding of dArc1 to candidate DNA sequences *in vitro* by detecting a mobility shift of the DNA-protein complex. Together, these three complementary approaches aim to determine whether dArc1 functions as a transcriptional regulator of its own gene. Understanding the mechanism of dArc1 autoregulation could provide new insights into Arc protein function and the gene regulatory pathways involved in neuronal function.

Are three-dimensional detonation waves in a confined domain "realistic"?

Sebastian A. Seda

Mentors: Guillaume Blanquart and Benjamin R. Breer

There are parameters that can affect the structure of a detonation wave such as boundary conditions and unburnt pockets of gas. However, a clear consensus on which parameters dominate remains a subject of ongoing debate. Detonation waves are inherently multi-dimensional due to intrinsic flow instabilities. 2D cases have been studied the most because computing 3D simulations using Direct

Numerical Simulation (DNS) with detailed chemistry in large domains is computationally expensive. Therefore, within our 3D simulations, DNS is performed with detailed chemistry to examine specific parameters such as boundary conditions, expanded domain size and finer grid resolution to determine how the detonation structure is altered. Four configurations are evaluated: (1) symmetrical boundary conditions in the y and periodic in z, (2) fully periodic boundary conditions in y and z, (3) an expanded domain size and (4) a refined grid resolution (both using fully periodic boundary conditions). The simulations provide detailed insights through single point statistics, probability density functions (PDF) of shock speed and induction length, planar averages, 3D isosurfaces of the shock front and the power spectral density (PSD) of temporally varying quantities. Overall, we aim to understand how boundary conditions, grid size and refined grid resolution affect detonation waves within a confined space.

Encoding and replicating hierarchical spatial information in reaction-diffusion systems

Eric F. Segrest

Mentors: Erik Winfree and Daichi Hayakawa

Living systems store information in geometry: when a cell divides, it copies not only its genetic sequence but also the spatial arrangement of its parts. Chemical mixtures that react and diffuse may spontaneously organize into spots, stripes, and other related patterns, some of which even reproduce themselves. However, no reaction-diffusion system has been shown to replicate a pattern with complex internal structure, in which distinct labeled regions are inherited by each copy. We ask whether such information can be copied by local chemical interactions alone, without templates or external control. Our target pattern, an American flag, is built hierarchically: a large chemical spot defines the boundary of each copy, smaller interior spots of two distinct types divide that boundary into labeled regions, and further reactions fill those regions with finer patterns. In simulation, asymmetric coupling between the two interior types favors one region of the first type and three of the second inside each dividing spot, breaking the symmetry that defeated earlier designs. Hierarchical coupling of simple chemical patterns may therefore preserve spatial organization across replication, suggesting a route toward chemically plausible self-replication of complex geometries.

On-chip mid-infrared optical parametric oscillation in thin-film lithium niobate-on-sapphire

Elizabeth Seward

Mentors: Alireza Marandi and Selina Zhou

Optical parametric oscillators (OPOs) based on second-order ($\chi^{(2)}$) nonlinear interactions provide a versatile route to coherent, widely tunable light at wavelengths inaccessible to direct laser gain media, making them a key tool for mid-infrared spectroscopy and sensing. While bulk and fiber-based OPOs have long served this role, their footprint, cost, and power requirements limit deployment outside the laboratory. The advent of thin-film lithium niobate (TFLN) nanophotonics, offering tight mode confinement and engineerable dispersion, has enabled chip-scale OPOs with dramatically reduced thresholds. However, on-chip OPOs have so far been largely confined to signal/idler pairs below $3.5\text{ }\mu\text{m}$, falling short of the spectral reach of their bulk counterparts.

This work presents an on-chip, doubly resonant OPO fabricated on a TFLN-on-sapphire platform, pumped by a pulsed telecom laser near 1550 nm . Using periodically poled, dispersion-engineered waveguides, we demonstrate parametric oscillation with emission spanning from degeneracy to $4\text{ }\mu\text{m}$. We characterize the device threshold behavior, measuring an oscillation threshold of 5.5 W on-chip pump peak power, and quantify cavity performance through finesse and free-spectral-range (FSR) measurements across the signal and idler resonances. On-chip signal and idler powers are calibrated using detector gain and filter transmission corrections, yielding a peak conversion efficiency of 0.4% . Pump-wavelength tuning is shown to continuously shift the idler emission, consistent with the resonator's group-velocity-matched design.

These results establish a compact, low-threshold source of tunable mid-infrared light directly accessible from a telecom pump, and point toward integrated platforms for on-chip molecular spectroscopy and sensing applications.

Designing a superconducting microwave resonator to magnetically couple to $^{171}\text{Yb}^{3+}:\text{CaWO}_4$

Nimisha Sivaraman

Mentors: Andrei Faraon and Chengyi Luo

Large ensembles of solid-state spins interacting through cavity-mediated interactions are important systems for both studying collective phenomena and developing quantum technologies, including quantum memory. One promising material platform for studying these collective dynamics is $^{171}\text{Yb}^{3+}:\text{CaWO}_4$ due to the combination of long intrinsic coherence and a strong microwave transition dipole moment. By coupling a bulk $^{171}\text{Yb}^{3+}:\text{CaWO}_4$ crystal to a 3D microwave resonator, the Faraon Group has observed extended Ramsey coherence time enabled by the photon-mediated interactions between ytterbium ions. Moving forward, a chip-scale platform with high quality factor microwave resonators strongly coupled to ion ensembles will enable scalable quantum memory. To this end, I have explored multiple 2D resonators using finite element softwares to specifically optimize coupling strength, mode filling factor, and magnetic field homogeneity. After converging on a single design that can be feasibly fabricated, the chip will be used for future experiments.

Lactadherin-modified biosensor for detection of inflammatory biomarker in perspiration

Jessica Slaughter

Mentors: Wei Gao and Jeonghee Yeom

The release of phosphatidylserine (PS)-exposed microparticles (MPs) during cell death and stress responses is a key biomarker of inflammation and has been associated with conditions such as severe decompression sickness (DCS), which disproportionately affects US Navy divers and individuals engaged in prolonged SCUBA diving. Current methods for measuring these biomarkers rely primarily on blood analysis, which is invasive and requires trained personnel, sterile sample handling, and laboratory processing, thereby limiting its suitability for continuous and in situ monitoring. Wearable electrochemical biosensors offer a promising alternative by enabling real-time, noninvasive detection of biomarkers in accessible biofluids such as sweat. Here, we present a wearable biosensor for the electrochemical detection of sweat MPs. The working electrode is functionalized with lactadherin, an affinity protein that selectively captures PS exposed on the outer leaflet of inflammatory MPs. To achieve sensitive detection, we employ a competitive binding assay in which methylene blue (MB)-loaded liposomes serve as signaling probes. As the concentration of PS-rich MPs increases, these MPs compete with MB-loaded liposomes for fixed binding sites on immobilized lactadherin, consequently decreasing the electrochemical signal. This platform has the potential to enable early identification of individuals at risk of severe DCS and continuous monitoring of inflammatory conditions in point-of-care and remote settings.

Developing improved computational algorithms for the stable Andrews–Curtis conjecture

Elliot D. Steffen

Mentor: Sergei G. Gukov

This project develops search algorithms for the stable Andrews–Curtis conjecture, which asks whether certain presentations of the trivial group can be simplified to a standard form using allowed transformations. Building on an existing substitution-based search, I implemented a change-of-variables move and developed Dual and Triple Greedy Search, which schedule stronger moves separately near search bottlenecks.

All three search configurations solved a benchmark of 640 previously solved presentations. The substitution-only baseline explored an average of approximately 4,963 states, compared with 338 for Dual Greedy Search. Triple achieved a similar reduction. Neither new algorithm solved any of the 124 unresolved presentations within a budget of one million explored states per case. Dual shortened 22 of them, while Triple shortened 28, including the 22 shortened by Dual. These results did not resolve the conjecture, but the new search configurations explored fewer states on the known-solvable benchmark and reached shorter forms in some unresolved cases.

Simulation-augmented learning-based control for a liquid payload robot

Sonal Sukhatme

Mentors: Soon-Jo Chung and Jose Leonardo Brenes Calderon

Liquid payloads on vehicles complicate autonomous operation by presenting increased risks and leading to more complex vehicle dynamics. Previous solutions to this problem use more conservative models, which may lead to limited performance. In this project, we utilize open-source CFD simulations of sloshing behavior to generate data for training learning-based controllers. Introducing simulation results provides more diverse training data, including scenarios that may be difficult to replicate in experimental conditions. Another challenge of estimating fluid dynamics is partial information of fluid state due to limited sensing. This project augments a newly developed composite adaptation control architecture that addresses these observability limitations. Both controller versatility and predictive capabilities can be improved due to the expanded training regime from simulations.

The fixed point property for countable groups

Elizabeth V. Teryoshin

Mentor: Omer Tamuz

A group G is said to have the fixed point property if every action of G on a compact convex set has a fixed point. For example, the infinite cyclic group $\mathbb{Z}$ has the fixed point property while the infinite dihedral group does not. Work of Oliver from 1975 classified which finite groups have the fixed point property, but the situation for infinite groups is still not known in general. In joint work with Nicolas Monod, we provide an example (that's not $\mathbb{Z}$) of an infinite, finitely generated group with the fixed point property.

Surface plasmon resonance spectroscopy for investigating the growth of molecular additive layers on copper for CO₂ reduction

Cayley Y. Tolbert-Schwartz

Mentors: Harry A. Atwater, Jr., and Ellis Spickermann

The electrochemical CO₂ reduction reaction (CRR) utilizes renewable energy sources and stores energy longer than any typical battery or capacitor, and can be done safely and cheaply in aqueous solutions. However, in these solutions, the CRR suffers from the uncontrolled occurrence of the hydrogen evolution reaction (HER), lowering the efficiency of CO₂ reduction. Molecular additives, particularly *n*-tolyl-pyridinium, have been previously shown to suppress HER. We found that in electrochemical cells, *n*-tolyl-pyridinium was deposited on the cathode when a voltage of around -0.5 V vs SHE or lower was applied, with the deposition rate increasing at lower voltages. When running the cell while sparging with Ar, the current approached 0 due to the lack of CRR and suppression of HER as the additive layer grew. These final additive layers were determined to be around 100 nm thick using ellipsometry. *In-situ* surface plasmon resonance (SPR) spectroscopy was also used to investigate *n*-tolyl-pyridinium layer growth rates. The redshift of the layer's characteristic SPR dip was examined at different growth times, and the layer thicknesses were measured using ellipsometry. The final thickness could not be examined as the SPR dip redshifted out of the range of the spectrometer. The interim thickness points were instead used to model the layer growth in COMSOL based on the change in the SPR spectrum across the whole growth process.

SPR analysis indicated that a copper-water interface reappeared when lower voltages were applied. Prior to any applied voltage, the SPR spectrum of simple copper-electrolyte showed an SPR dip at around 570 nm. At an applied voltage of around -0.5V vs SHE the additive layer grows and the SPR dip redshifts past 900 nm, out of the range of the spectrometer. Regardless of the thickness of the additive deposited beforehand, at around -1.2 V vs SHE, an SPR dip at 575 nm appeared and the

general shape of the spectrum became similar to that of a simple copper-water interface. Below -1.2 V, the SPR dip redshifted 10-20 nm. COMSOL modeling indicated that these results could be generated by the presence of a water layer which decreased in thickness as the voltage decreased. This water layer could be formed by the additive momentarily peeling off of the copper due to swelling or bubbling from gas production. This water layer disappeared once voltages were returned to above -1.2 V.

Utilizing these observations, we aim to control the thickness and growth rates of the *n*-tolyl-pyridinium layer in CO₂ reduction cells. We plan to control the final layer thickness by changing the concentration of *n*-tolyl-pyridinium in the electrolyte, and control the final growth rate by changing the applied voltage.

Adapting pharmaceuticals for *in vivo* manipulations of rove beetle chemoreceptors

Y Lam Voong

Mentors: Joseph Parker and Hayley Smihula

Manipulating chemoreceptors via pharmaceutical methods can enhance studies on chemosensory contributions to animal behavior. Here, we aim to adapt pharmaceuticals for *in vivo* use in rove beetles, a family of interest due to many species having evolved socially parasitic relationships with ants. Our goal is to establish protocols for treating adult beetles and manipulating TRPA1 and Orco chemoreceptors in *Dalotia coriaria*. We study HC-030031 and VUAA1, compounds that target TRPA1 and Orco, respectively. These chemicals can be delivered via injection, volatilized exposure, topical and oral application. We performed a series of titrations to determine the necessary combination of delivery method, concentration, and persistence time that can disrupt receptor functioning during behavioral assays. Previously established TRPA1 *Drosophila melanogaster* and Orco *Dalotia* mutants were used to evaluate the efficacy of behavior assays for determining successful HC-030031 and VUAA1 treatments. We improved delivery parameters to prevent toxicity that can interrupt behaviors in biological assays. In the future, we hope to translate these protocols to the myrmecophilous rove beetle *Sceptronius lativentris*, which is less genetically tractable due to its lifestyle living inside of ant colonies. Our research on this species would therefore benefit from having a non-genetic method for disrupting chemoreceptor function.

Engineering explosive synchronization in an 8-node NEMS oscillator network

Alice T. Wang

Mentors: Michael L. Roukes and Scott T. Habermehl

Explosive synchronization, which occurs when all nodes in a coupled oscillator network synchronize suddenly at a critical coupling strength, is theorized to be an important mechanism for neural communication. We aim to demonstrate explosive synchronization experimentally using an 8-node NEMS network. We first analyze phase transitions produced by uniform frequency distributions in numerical simulations and experiments on the NEMS network, and find that hysteresis reflects the presence of multi-stable states, with the presence of "cluster states" occurring as a finite size effect. Using numerical simulations, we further determine that explosive synchronization could be produced by frequency-weighting of coupling strength or correlating intrinsic frequencies with coupling topologies in sparse Erdos-Renyi random graphs. We find that although explosive synchronization is guaranteed only in the continuum limit, it could be implemented in a small-node experiment using specific network topologies in the low frequency limit. Our results represent a step forward towards improved analog reservoir computing with NEMS.

Neural circuit mapping and behavioral characterization in activity-based anorexia and restraint stress models

Ava Arabshahi

Mentors: Yuki Oka and Tongtong Wang

Stress can suppress appetite through multiple neural pathways, but whether distinct stressors recruit shared or separate brain circuits remains unclear. This project compares restraint stress (both acute and chronic immobilization) with activity-based anorexia (ABA), an established model of stress-induced hypophagia used here as a positive control. In ABA, time-restricted food access paired with a running wheel drives paradoxical hyperactivity and rapid weight loss. Whole-brain Fos immunofluorescence is used to map neural activation across the two models and their respective controls. Mice were assigned to five groups (ABA, food restriction, wheel access, restraint, and control). Wheel-running revolutions were tracked using a self-built Arduino logger. Additionally, body weight, food intake, and water intake were recorded daily for 14 days. ABA mice lost approximately 19% of baseline body weight at their lowest point, compared with about 6% in restraint-stressed mice and negligible change in controls. Restraint-stressed mice maintained near-normal, slightly increased food intake despite this weight loss, indicating that their weight loss may arise through a mechanism separate from ABA's intake-driven loss. Brain sections are being processed to quantify Fos-positive cell density across hypothalamic, amygdalar, and brainstem regions to determine whether the two models engage dissociable pathways.

Analysis of the exploratory dynamics of microtubule search for kinetochores during cell division

Nishanth Chavourkar

Mentors: Rob B. Phillips and Minakshi Ashok

Many biological processes are designed to reach specific goals to ensure proper functioning of an organism. During cell division, microtubules emerging from the poles of a spindle must each reach individual chromosomes. Any error in this process can lead to one of the daughter cells having one too many or one too few chromosomes. Thus, biological systems must develop mechanisms for reaching their goal with high fidelity. To achieve high fidelity, many biological systems operate on a paradigm of variation and selection, where a biological process explores multiple trajectories, terminating unsuccessful outcomes until a successful outcome is reached. We aim to provide a more realistic description of the variation and selection paradigm within the context of the exploratory dynamics of the mitotic spindle. To achieve this we accounted for the dynamics of the termination of a trajectory (microtubule catastrophe), and analyzed how the structure of the mean search time for a chromosome varied as we introduced this correction. This work is part of a broader effort to describe the dynamics of biological systems using the language of trajectories. An understanding of the trajectories taken by biological systems allows us to further understand *how* a biological system reaches a specific goal.

Real-time attentional classification using smart glasses

Nia P. Desai

Mentors: Wei Gao and Ruixiao Liu

Attentional states are the result of coordinated neural, cardiovascular, and ocular activity. However, these signals are typically analyzed independently without consideration of relationships between physiological systems. This project uses a multimodal smart-glasses platform integrating electroencephalography, photoplethysmography, and eye-tracking data to quantify and characterize attentional states. We implemented three attention tasks used to induce various attentional states and collected data from six participants. Signal-processing pipelines were developed to extract task performance and features from neural, cardiovascular, and eye image data. For the eye-tracking images, a custom landmark-based model was used to simultaneously extract pupil, blink, and eye-state features. These three data streams have been integrated into a machine-learning model to analyze features across modalities that are not apparent when analyzing signals independently. By

comparing unimodal and multimodal relationships, this work aims to enable closed-loop, real-time classification of attention from interactions between neural, cardiovascular, and ocular data into a single wearable system. This work establishes a foundation for wearable systems capable of continuous attention monitoring, with potential applications in driving safety, workplace performance, healthy aging, and neurological disease.

Synthesis and characterization of a trimetallic MNi₂ complex for CO₂ reduction

Anna L. Liu

Mentors: Theodor Agapie and Nicole Piekut

Multimetallic catalysts enable enhanced or divergent reactivity from traditional monometallic catalysts owing to cooperative interactions between metal centers, with applications in small molecule activation and organic transformations. Here, we report progress towards synthesizing a trimetallic MNi₂ complex (M = Ti, Zr, Mo) stabilized by a rigid, planar, tetradentate bis-triazinyl bipyridine (BTBP) ligand. This system integrates late-late and early-late transition metal interactions for merged advantages including multiple metal-metal interactions to tune reactivity, multiple redox active centers to store electrons, and multiple metal sites for substrate binding. As such, we anticipate this unique structure to be well suited for CO₂ reduction. The BTBP ligand was synthesized and metalated with various titanium sources; precursors containing other metals (M = Zr, Mo) were also investigated. Nuclear magnetic resonance (NMR) characterization suggests that TiCl₃ and Ti(NtBu)Cl₂(py)₃ successfully metalated the BTBP ligand, and further single-crystal X-ray diffraction (XRD) studies will be conducted to confirm the structures. Subsequently, nickel will be added to the M-BTBP complexes, screening different conditions and nickel sources. The geometric and electronic structure of the final MNi₂ complex will be probed with single-crystal XRD, NMR, X-ray absorption spectroscopy, and computational studies. Finally, the aptitude of these complexes for CO₂ reduction will be studied.

Investigating the role of insulin signaling in zebrafish meningeal and neural development

Delia N. Parco

Mentors: Marianne Bronner and Ayyappa Raja

Insulin-like growth factor (IGF) signaling is an evolutionarily conserved pathway crucial for animal growth and development. Previous research has demonstrated that IGF binding protein 2 (IGFBP2) plays a major role in brain development and tumorigenesis; however, the specific role it plays in cell type development is not yet known. Using zebrafish (*Danio rerio*) as a vertebrate model, the lab generated a genetic knockout of *igfbp2a*, a zebrafish homolog of human IGFBP2, using CRISPR-Cas9 genome editing to investigate its role in embryonic development. 7-base-pair deletion mutants exhibit optic tectum closure defects by 48 hours post-fertilization, with phenotypic severity distributed in a pattern consistent with expected inheritance. In situ hybridization chain reaction experiments reveal a dramatic decrease in *igfbp2a* expression in the meninges of homozygous mutants, accompanied by a decrease in differentiated neuron markers and an increase in neural stem cell (NSC) markers, indicating self-cycling of NSCs. These findings suggest that *igfbp2a* may regulate brain development by modulating the brain meninges axis. Future work will be done in the lab to better understand the role of *igfbp2a* in insulin signaling during brain development. Together, these results establish zebrafish *igfbp2a* mutants as a viable model for investigating IGF signaling during development and neurogenesis.

Establishing genetic methods for chemoreceptor manipulation in rove beetles

Conner I. Sakae

Mentors: Joseph Parker and Hayley Smihula

Pasadena, California is home to *Sceptrobius lativentris*, a rove beetle species and obligate symbiont of the ant *Liometopum occidentale*. *Sceptrobius*' lifestyle stands in stark contrast to most rove beetles, such as *Dalotia coriaria*, free-living generalists that are prey to ants. We wish to understand how chemosensory systems for ant detection vary between these two beetles. By applying genetic manipulations with RNAi and CRISPR, we disrupted the function of two chemoreceptors involved in pheromone sensing, TRPA1 and Orco, in *Dalotia* and *Sceptrobius*. We created and tested RNAi constructs to knockdown TRPA1 in both species, utilized an existing *Dalotia*-Orco CRISPR mutant and

have begun validations on a *Dalotia*-TRPA1 CRISPR line. High throughput assays were developed to assess receptor disruption which were subsequently validated with more comprehensive behavioral analysis. Control odors were screened with the goal of producing a consistent behavioral assay that could evaluate the success of genetic manipulations on the phenotypic level. With genetic tools and control assays that can successfully alter and assess chemoreceptor function, we can apply these methods to future work regarding the detection and behavioral response to ant pheromone compounds as a part of ongoing research into the evolution of *Sceptobius*' myrmecophilic lifestyle.

Neural and behavioral correlates of social connection during human-AI interaction

Anika Shethia

Mentors: Dean Mobbs and Emily Towner

Social disconnection is a growing global concern, and people are increasingly turning to artificial intelligence (AI) for social support. Yet the neural and psychological mechanisms through which social connection forms during naturalistic social interactions, whether with humans or AI, remain poorly understood. This project applies switching linear dynamical systems (SLDS), a mathematical framework for modeling how internal brain states evolve over time, to test whether discrete states of connection and disconnection can be recovered empirically from neural data. We further ask how social connection alters the way we cognitively represent others, adaptively mentalize about them, and respond to their influence. Thirty participants hold text conversations with artificial partners designed to foster connection or disconnection while undergoing functional magnetic resonance imaging (fMRI). Participants also complete a face-viewing task before and after conversation, a Rock-Paper-Scissors task, and an artwork-rating task, indexing self-other neural overlap, adaptive mentalizing, and susceptibility to social influence respectively. Pilot behavioral data from a separate online sample (N = 40) confirm that the manipulation produces the intended difference in felt connection, validating the paradigm for the scanner phase. Scanning is ongoing; analysis will test whether latent connection states predict mentalizing depth and influence.

Developing a human tonsil organoid workflow for single-cell antibody discovery using the Beacon optofluidics platform

Arman M. Thoresen

Mentors: Pamela J. Bjorkman and Kate Malecek

Human tonsil organoids recapitulate key features of human adaptive immunity and provide an *ex vivo* model for studying vaccine responses. The goal of this project is to develop a workflow for isolating and analyzing vaccine-induced plasmablasts and plasma cells (PB/PCs) from tonsil organoids using the Beacon, a single-cell optofluidics platform that measures antibody secretion and enables recovery of individual cells for antibody characterization. Tonsil organoids were stimulated with live attenuated influenza vaccine (LAIV) or influenza hemagglutinin nanoparticles, while unstimulated organoids served as controls. Flow cytometry was used to identify antigen-specific B cells and PB/PCs, and magnetic cell separation was used to enrich PB/PCs for Beacon analysis. Initial experiments demonstrated antigen-specific B-cell and PB/PC responses following LAIV stimulation, while unstimulated organoids showed minimal antigen-specific responses. Hemagglutinin nanoparticles also induced antigen-specific responses, although generally at lower levels than LAIV. Magnetic cell separation enriched PB/PC populations, supporting its use for preparing antibody-secreting cells for Beacon analysis. These results establish a foundation for optimizing PB/PC isolation and antibody secretion from tonsil organoids, enabling future single-cell antibody discovery and characterization of responses to next-generation vaccine candidates., enabling future single-cell antibody discovery and characterization of responses to next-generation vaccine candidates.

Towards engineering *Azotobacter vinelandii* for characterization of non-canonical nitrogenase

Hanna J. Timek

Mentors: Smruthi Karthikeyan and Zahra Shivji

Current industrial methods of ammonium synthesis are energy-intensive and unsustainable. Hence, the development of bioengineered nitrogen fixation systems is urgently required to sustainably meet global nitrogen demands. A possible solution is engineering nitrogenase enzymes in order to maximize

their ammonium output for commercial use. Although the enzyme is found in an immense variety of organisms, current knowledge of nitrogenase comes almost exclusively from the aerobic, soil-dwelling Gammaproteobacteria *Azotobacter vinelandii*. The limited capacity to study nitrogenases outside of *Azotobacter vinelandii*, a culturable diazotroph, creates significant gaps in understanding their diversity, functionalities, and characteristics. In this project, we use CRISPR-associated-transposon machinery to develop a platform to study non-culturable nitrogenases in the model organism *A. vinelandii*. We engineered *A. vinelandii* to be non-diazotrophic by removing the *nifHDKEN* cluster of its nitrogenase gene via homologous recombination. Currently, we are optimizing our protocol for replacing the removed cluster by using a plasmid harboring both the *nifHDKEN* cargo and CAST module. In the future, we hope to use the system to insert cargo of non-canonical nitrogenases and study their capacity to fix nitrogen using acetylene reduction assays and ^{15}N isotope labeling.

CellSAM3: A promptable foundation model for cell segmentation and tracking

Edmund Tsou

Mentors: David A. Van Valen and Ross Barnowski

Cells are a fundamental unit of biological organization, and most quantitative cellular experiments begin with cell segmentation to identify individual cells in imaging data. Accurate segmentation is essential because morphology, expression, and dynamics are measured per cell, so errors propagate downstream. Although deep learning has enabled generalist models that segment cells accurately across imaging domains, these models provide limited support for user-guided correction at inference time. Here, we present CellSAM3, a promptable foundation model for cell segmentation. CellSAM3 builds on the Segment Anything Model 3 (SAM 3) by jointly fine-tuning a shared backbone for automatic cell detection and promptable segmentation using 49 public datasets spanning fluorescence microscopy, phase-contrast microscopy, and histopathology. We show that a single model can automatically segment all cells in a field of view, respond to point and box prompts, and produce nuclear or whole-cell masks as specified by a text prompt. CellSAM3 outperforms all evaluated generalist models on an internal benchmark and matches the leading method on the held-out NeurIPS 2022 Cell Segmentation Challenge. We adapt CellSAM3 to time-lapse imaging using SAM 3's memory-based video tracker, which preserves cell identities across frames. These results establish CellSAM3 as a unified framework for automatic, interactive, and temporal segmentation.

Characterizing and reducing earthquake-related control noise at LIGO Hanford

Matilda G. Damon

Mentor: Elenna Capote

Seismic ground motion can cause the Laser Interferometer Gravitational-Wave Observatory (LIGO) detector at Hanford to exit its operational state. This project investigates whether subtracting noise associated with interferometer control signals during Earthquake Mode, a control configuration used during periods of elevated seismic motion, can improve the detector's binary neutron star (BNS) detection range. Earthquake periods were compared with nearby intervals of normal operation using measurements of signal amplitude, frequency content, and coherence between control signals and the gravitational-wave strain channel. Transfer-function-based subtraction pipelines were developed and evaluated by comparing the BNS detection range before and after subtraction. Preliminary results show that linear subtraction of selected longitudinal control signals recovered approximately 10 megaparsecs of BNS range. The NonSENS noise-subtraction framework was also implemented and evaluated for linear and nonlinear noise subtraction that is not well described by simple transfer functions. In parallel, control signals are being compared with earthquake-monitoring channels and detector state information to identify which systems require increased control authority during earthquake mode and preceding lock loss. These results indicate that targeted control-noise subtraction can improve detector sensitivity during earthquake mode, enabling increased observing time.

Characterizing quasi-periodic low-frequency noise during O4 at LIGO Livingston

Melanie H. DeMarcay

Mentor: Jane Glanzer

During LIGO's fourth observing run, the Livingston detector's binary neutron star (BNS) range exhibited recurring oscillations, appearing as noise in the 30-50 Hz band of the gravitational-wave strain signal. This project uses frequency-domain analysis to identify auxiliary channels that may be responsible for this noise with the goal of mitigating it before the next observing run. It is important to reduce this excess noise as it is located in the region where the interferometer is particularly sensitive to the mergers of high mass binary black holes. We investigated whether auxiliary channels were coupled with the strain oscillations by using coherence and bicoherence analyses. The survey comparing ordinary coherence between the strain and thousands of auxiliary channels found no channel with a consistent linear relationship, suggesting the coupling may instead be nonlinear. We have since shifted to a segmented bicoherence survey, sensitive to quadratic phase coupling between frequency components, which has revealed structured, localized patterns for several channels within the oscillation's frequency band. Continued analysis across additional detector subsystems and oscillation days is needed to determine whether these patterns indicate a genuine physical driver of the noise.

Quantum noise lock optimization for tabletop Waveguided Optical Parametric Amplification (WOPA)

Abigail S. Gebremariam

Mentors: Rana X. Adhikari and Peter Carney

The Laser Interferometer Gravitational-Wave Observatory (LIGO) detectors' sensitivity is affected by various noise sources, including quantum noise. LIGO currently reduces this noise using frequency dependent squeezing generated by an Optical Parametric Oscillator (OPO) with active cavity locking. To simplify the setup, this project uses a waveguide Periodically Poled Lithium Niobate (PPLN) crystal to generate squeezing without requiring cavity locking. To observe squeezing, the squeezed vacuum is combined with a Local Oscillator (LO) at a 50/50 beam splitter and detected by two photodetectors in a Balanced Homodyne Detection (BHD) configuration. By using a Piezoelectric Transducer (PZT)-actuated mirror, the LO phase is adjusted to measure both squeezed and anti-squeezed quadratures. Quantum Noise Locking (QNL) stabilizes the relative phase between the LO and squeezed vacuum to

enable an accurate squeezing measurement. The objective of this project is to optimize the QNL feedback loop to minimize residual phase noise. This is achieved by generating stronger error signals, applying control loop filters, and calibrating overall loop gain for stable feedback performance. Additional measurements were conducted for this experiment, such as photodetector noise and response characterization, along with design modeling of optimal photodetector dynamics for future squeezing measurements in LIGO's audio detection band.

Measuring environmental noise for Cosmic Explorer site selection

Shreya Goel

Mentors: Robert Schofield and Michael R. Landry

Cosmic Explorer (CE) is a proposed next-generation gravitational wave observatory to be built in the US, consisting of an observatory 10x the size of current LIGO observatories. Working in concert with the European Einstein Telescope, CE will observe gravitational wave sources across cosmic time, including black hole mergers from the first stars to form in the universe. To select a site for CE, it is important to understand sources of potential noise from various sources, including seismic, electromagnetic, and more. We are working on improving standalone datalogger kits, consisting of seismometers, magnetometers, and microphones, used to assess potential sites. We have developed a new kit that uses two seismometers in tandem to understand surface wave propagation from various sources. We have also written new code that uses this data to calculate wave velocity, which is useful to understand the quality factor of different sediment types and interpret coupling. We tested our kit at various locations around LHO to better understand local wave propagation velocity, as well as at Goose Prairie Solar Farm in order to understand the constraints of solar farm locations near gravitational wave observatories. Results from this study will be helpful in future site evaluation for CE.

Cavity-enhanced sum frequency generation for 2 μm interferometer readout

Rohit V. Jain

Mentors: Rana X. Adhikari and Radhika Bhatt

Next-generation gravitational-wave detectors are exploring interferometry at 2 μm laser wavelength to reduce thermal noise sources and improve detector performance. Efficient photodetection at these wavelengths remains a challenge due to the lack of high-quantum-efficiency detectors. Sum-frequency generation (SFG) offers an alternative way by upconverting 2 μm light to visible wavelengths, where high-QE detectors are readily available. In the proposed experiment, a strong 1064 nm pump beam is resonantly enhanced in a bow-tie cavity and combined with a weak 2 μm signal beam inside a PPLN crystal to produce 700 nm light. This project works toward a proof-of-principle demonstration of this scheme, beginning with locking and characterizing the 1064 nm cavity before introducing the crystal and signal beam. The experimental work focused on cavity alignment, beam profiling, mode matching, and characterization of the 1064 nm pump laser, as well as PDH locking of the cavity. These developments in this work revealed limitations in the existing optical system and provided a better understanding of the requirements for stable operation, bringing the cavity closer to a stable lock and laying the groundwork for crystal integration and SFG measurement planned in subsequent stages of the project.

Efficient modeling of spin evolutions in precessing neutron star binaries

Quinn A. Kasdan-Grollo

Mentors: Giulia Fumagalli and Taylor Knapp

Spin-induced precession in a compact binary has a significant impact on both the detected gravitational wave signal and the binary's evolution history, making modeling precessing compact binaries an interesting avenue for study. We extend upon the post-Newtonian framework developed by Gerosa *et al.* 2023, which modeled the evolution of precessing black hole binaries, to the unexplored regime of binary neutron stars (BNSs) by introducing precession-averaged equations generalized to BNSs in quasi-circular orbits. This generalization makes finding the bounds for which angular momentum vectors are allowed to precess between non-analytical and less robust compared to the binary black hole case. To resolve this, we identify restrictive physical bounds on the

interpolation grid used to find the aforementioned precession bounds, further optimize the interpolation grid, and compactify spin-precession equations making evolution of precessing BNSs robust and efficient. This new framework for modeling BNSs efficiently tracks the evolution of spin directions throughout the entire inspiral, enabling future study of the impact of matter effects on spin precession. Understanding this impact has the potential for constraining neutron star equations of state through study of real event spin evolutions, especially in light of next-generation gravitational wave detectors where the long BNS inspiral is observable.

A system identification framework with RL-based training for technical noise reduction in the Caltech 40m Prototype

Michael N. Padilla

Mentors: Nathan A. Holland, Christopher Wipf, and Rana X. Adhikari

To achieve the sensitivity targets of future gravitational wave detector upgrades, low-frequency technical control noise must be significantly reduced. We propose the development of an accurate digital twin model to serve as the environment for fast offline training of a reinforcement learning-based control policy that can be implemented directly into the real system. This presentation details our development of a physics-based digital twin of the Caltech 40m Prototype Input Mode Cleaner – a triangular, suspended optic cavity. We implemented a gray-box model directly within LIGO's real-time control architecture, constructing the complete signal path to accurately capture the dynamics of the digital controller and electronics alongside the physical suspended optics. We also demonstrate the implementation of an optimal system identification algorithm for measurement of the 40m system. This algorithm allows for fast and optimal measurements that will serve as the basis for further validation and fitting of the digital twin, enabling a sufficiently robust training environment for effective control policy implementation onto the real system.

Improving mode-matching and BHD visibility for waveguide-based squeezing at 1064 nm

Nishant G. Pawar

Mentors: Rana X. Adhikari and Peter Carney

Waveguide-based Optical Parametric Amplification (WOPA) is a promising alternative that could simplify the conventional cavity-based squeezed-light sources used in the LIGO detectors: a single pass of 532 nm pump light through a periodically poled lithium niobate (PPLN) waveguide produces squeezed vacuum at 1064 nm without any cavity that has to be locked. How much of this squeezing we actually detect, however, depends strongly on the visibility of the Balanced Homodyne Detection (BHD), i.e. on how well the Local Oscillator (LO) beam is mode matched to the squeezed field. This project aimed to improve the BHD visibility of the WOPA tabletop setup. Because squeezed vacuum carries no classical power, we used a bright 1064 nm beam sent through the waveguide as a stand-in for the squeezed mode, profiled both it and the LO, and designed a new two-lens mode-matching solution placed before the PZT phase-scanning mirror. With this solution the visibility improved from 90% to 91.2%, and the detected squeezing increased from -3.0 dB to -3.20dB (with +3.3 dB of anti-squeezing) at 142 mW of pump power. Using a loss model calibrated to this measurement, we predict that reaching 95 - 97% visibility would raise the detected squeezing to about -3.6 to -3.9 dB at the same pump power, and to about -4.5 to -4.9 dB at 250 mW. A Gaussian-beam (q-parameter) analysis shows that the lens solution by itself allows a visibility close to 99%, so the remaining gap is dominated by alignment and by the non-ideal, elliptical beam leaving the waveguide.

A physics-accurate digital twin and system identification framework for the Caltech 40 m prototype input mode cleaner

Arush Pimpalkar

Mentors: Rana X. Adhikari and Nathan A. Holland

Gravitational-wave observatories such as LIGO depend on active control systems to hold their suspended mirrors aligned and on resonance, and at low frequencies the noise these controllers inject onto the mirrors sets the instrument's astrophysical reach. Developing and testing better controllers on the real hardware is slow and risky, so this project builds a "digital twin", a validated software replica, of the Input Mode Cleaner, a three-mirror suspended optical cavity at the Caltech 40 m

prototype interferometer. The twin has three parts: a mechanical model of how ground vibration reaches each mirror, an optical model of how mirror motion appears in the cavity's sensors, and an integration stage that links the two through the digital control loops and tunes their free parameters to match the real instrument. Each component was checked piece by piece against measurements taken directly on the hardware, including the mirror-damping loops and the alignment sensors, and reproduces the measured behavior well near the suspension resonances. The validated twin provides a safe, cost-free environment for designing and stress-testing control strategies, intended as a training ground for machine-learning controllers that could later transfer to the full-scale detectors.

Progress toward pulse-locking an optical enhancement cavity for pulsed laser amplification at Blue Laser Fusion

Jack M. Ryan

Mentors: Rana X. Adhikari and Georgia Bolingbroke

Blue Laser Fusion (BLF) is working to establish the world's first commercial inertial fusion energy, building on the National Ignition Facility's demonstration of fusion target gain in 2022. BLF utilizes Coherent Beam Combining (CBC) technology and an Optical Enhancement Cavity (OEC) to produce an efficient high-energy pulsed laser system. Because fusion ignition requires extremely high energy densities, laser fusion systems must operate in a pulsed regime to maximize peak power on target. Prior work by the Adhikari group at Caltech with a 1.5m prototype cavity operating in the continuous wave regime achieved an enhancement of 58,461, yielding kilowatt-level circulating powers. This project extends that work towards pulsed operation. The primary objective is to demonstrate pulse-locking of the laser to the OEC, characterize the pulse bandwidth and temporal shape, and understand the technical challenges involved in maintaining a pulse lock. Work to date has focused on reestablishing the past setup within an improved vacuum system. This included reinstalling 3.4 ppm loss mirrors, profiling and aligning the beam to the cavity axis, and achieving a lock under CW operation utilizing the Pound-Drever-Hall technique. Challenges arising from low optical coupling efficiency to the OEC were also explored, with several solutions being investigated.

Data acquisition, signal processing, and noise characterization for the GQuEST experiment

Ciara P. Treanor

Mentors: Lee P. McCuller, Alexander Ramirez, and Sander Vermeulen

The Laser Interferometer Gravitational-Wave Observatory (LIGO) has successfully detected gravitational waves from astronomical events. Previously developed theory suggests that a similar Michelson interferometer design can be used to detect small, stochastic spacetime fluctuations induced by entropy ("geontropic" fluctuations). This is the motivation for the Gravity from Quantum Entanglement of Space-Time (GQuEST) experiment. In order to perform this experiment, we need to characterize the noise of the main hardware systems. GQuEST will utilize a single-photon detector to detect signals below the quantum shot noise threshold, but it also uses a separate homodyne readout scheme which measures cross-correlated signals from two photodetectors. The experiment uses a 1550 nm laser and an analog-to-digital converter (ADC) to read voltages from modified photodetectors. The laser, the ADC, and the modified photodetectors each have some inherent noise, which we need to quantify before performing the experiment in order to distinguish real signal caused by geontropic spacetime fluctuations from background electronics and laser noise. In this report, we characterize these noise sources and establish a preliminary estimate of GQuEST's data acquisition sensitivity, based on theoretical and early experimental noise floor measurements.

Optimization of alignment sensing and control loops for the Jitter Attenuation Cavity (JAC)

Tran Bao Khanh Vu

Mentor: Jennifer L. Wright

Small angular misalignments in the laser beam entering the LIGO Hanford interferometer introduce noise that limits its sensitivity in the 100–1000 Hz range, an important frequency band for gravitational-wave detection. To reduce this noise, the Jitter Attenuation Cavity (JAC) was installed to spatially filter the beam and must suppress beam motion by 30 dB to meet the target for the next observing run, O5. We characterize the JAC's jitter suppression by commissioning and evaluating its

alignment sensing and control system through simulations and measurements. We also optimize the alignment and stability of the cavity control loops and troubleshoot issues that arise during commissioning.

Constraining astrophysical parameter bias due to glitches in future ground-based gravitational-wave detectors

Samantha F. Walker

Mentors: Sophie Bini and Jake S. Summers

Transient detector noise artifacts, or “glitches,” that coincide in time and frequency with gravitational-wave signals can bias the inferred astrophysical parameters of compact binary coalescences, particularly the black hole spins. We investigate how this bias evolves as detector sensitivity increases by simulating binary black hole signals based on GW200129 and GW191109, contaminated by glitches, under projected sensitivities for LIGO A# and Cosmic Explorer (CE). We consider two limiting scenarios in which either the glitch amplitude or the glitch signal-to-noise ratio (SNR) is held fixed relative to present-day detectors, and perform Bayesian parameter estimation on each simulated data set. We find that glitches injected with fixed SNR produce progressively smaller absolute biases in the effective spin parameters as the signal SNR increases, whereas glitches injected with fixed amplitude continue to induce biases of comparable magnitude despite substantially narrower posterior distributions. In the case we test, lowering the minimum analysis frequency from 20 to 14 Hz further reduces the fixed-SNR glitch bias to a negligible level. Adopting a signal SNR of 150 as the threshold above which fixed-SNR glitch biases become negligible, we simulate a population of binary black hole mergers and find that only 0.03% of detections in an A# network and 2-5% in a CE network exceed it, suggesting that most future detections would still require mitigation if they overlapped a comparable glitch. These results show that the impact of glitches in future detectors depends critically on how glitch populations scale with detector sensitivity, and provide insight into when computationally expensive glitch-mitigation methods remain necessary.

Cavity-enhanced frequency upconversion of 2 μm light for quantum-efficient detection

Kato J. Watsen

Mentors: Radhika Bhatt and Rana X. Adhikari

Proposed cryogenic upgrades to LIGO would use 2 μm light for low absorption in crystalline silicon test masses. However, photodetectors at 2 μm have quantum efficiencies below the near-unity levels achievable with photodetectors at shorter wavelengths. Sum frequency generation (SFG) of the 2 μm signal with a 1064 nm pump produces a 700 nm field, enabling the use of silicon photodetectors with quantum efficiency greater than 99%. Cavity-enhanced SFG improves the conversion efficiency by utilizing the high intracavity pump intensity. Here, the 1064 nm pump will be locked to a bow-tie cavity while the 2 μm signal field single-passes through the cavity, undergoing SFG with the pump in periodically poled lithium niobate. The SFG conversion efficiency will be experimentally measured and compared to analytical expectations.

Flattening and Construction of the GQuEST End Mirrors

Matthew C. Winkowski

Mentors: Lee P. McCuller and Daniel Grass

The GQuEST experiment searches for high frequency quantum gravity fluctuations. There are a number of issues that makes detecting these “geontropic” fluctuations challenging. The first of those, thermal noise, is directly proportional to mirror thickness, so a key noise reduction method for many experiments is using thinner mirrors. A thinner mirror can reduce the magnitude of one noise, but are more susceptible to curving under coating stress. This curvature can lead to deflections and scattering of the incoming beam, causing optical losses and experimental error. In order to correct for these losses I implemented a flattening method using a custom mount. I raster-scanned the mirror, in both

a flattened and unflattened state, analyzing coupling between the input and output modes. After assembling a second mount and mirror, they were put together into a test interferometer to further develop the angular and positional control loops. The control loops were then tested and verified for future implementation in the full-sized GQuEST experimental setup, improving the accuracy in the final search for quantum gravity.

Active magnetic field stabilization for cold-atom precision measurement

Martin Zanazzi

Mentors: Lee P. McCuller and Jeffrey D. Wack

The Rubidium Quantum Sensing (RbQ) experiment transfers optical signals into rubidium-87 hyperfine populations for temporally selective single-photon detection, in eventual support of GQuEST. Rb-87 hyperfine coherence is sensitive to ambient magnetic field noise, which induces Zeeman-effect decoherence. Because conventional mu-metal shielding is difficult to reconcile with the experiment's vacuum, space, and beam-access constraints, this project develops an active magnetic field stabilization system targeting the lab's dominant 60 Hz noise and harmonics. A compact Helmholtz-coil apparatus was built to validate the sensing and feedback architecture, suppressing uniform and low-order spatial field variations down to a sensor-limited noise floor near 4 nT/ $\sqrt{\text{Hz}}$ ahead of a future vacuum-compatible implementation. Work began by characterizing and comparing candidate magnetic sensors and readout electronics, including an anisotropic magnetoresistive sensor that requires periodic high-current reset pulses and a tunneling magnetoresistive sensor with a modestly higher noise floor but no reset requirement. Several instrumentation amplifiers and voltage regulators were also tested in parallel, with the design ultimately settling on the tunneling magnetoresistive sensor for its simplicity. These findings informed the design, components, simulation, and fabrication of three custom PCBs (sensor, signal-conditioning, and coil-driver boards) implementing the full feedback loop.