STRATA: An agentic training architecture for scientific task automation
Nicholas D. Abram
Mentors: 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.

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 understand-ing the processes that underlay their formation. Of particular interest are stripped-envelope core-collapse supernovae (SESNe), which have their hydro-gen 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.

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.

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 measure-ments 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 super-fluid 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.

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.

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.

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.

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 N_2, 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 N_2 adsorption by -0.40 eV, while H adsorption is unfavorable by +0.71 eV, making the site a leader in selectivity. Although Mo3 binds N_2 more strongly (-1.36 eV) and shows that the N_2 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.

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 a recently expanding class of short-lived X-ray bursts whose origins remain poorly understood. While recent advancements in soft X-ray surveys have enabled the detection of and rapid follow-up on several FXTs in the soft X-ray range, leading to the identification of various multi-wavelength counterparts, the nature of FXTs found in archival hard X-ray observations remains largely unexplored. In this work, we conduct a systematic multi-wavelength search for the counterparts to seven FXTs found in archival data taken by NuSTAR, a mission that observes the hard X-ray sky in the 3-79 keV range. Due to the limited localization accuracy of NuSTAR and the absence of contemporaneous follow-up, identifying genuine counterparts remains challenging. We search for associated emissions in the optical, radio, and ultraviolet using archival survey data and aim to provide constraints on the presence of counterparts to hard X-ray FXTs in these regimes.

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.

Diffusion guided posterior sampling for Bayesian inverse problems
Nicholas Z. Bai
Mentors: Anima Anandkumar and Jiachen Yao

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.

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 (SmI_2) is a versatile single-electron reductant widely employed in organic synthesis. Although SmI_2-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 CO_2 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 CO_2 concentrations vary with climate, reconstructed paleoatmospheric CO_2 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 CO_2. 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 CO_2 reconstructions to better determine whether this flora represents glacial or interglacial conditions and improve our understanding of Amazonian vegetation responses to past climate change.

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 cm^2 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.

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* d*ianthus, 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.

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 mKapα2 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 Mars-analog drone landing
Juan E. Beron
Mentors: Aaron D. Ames, Joe Moeller, Pedro Roque, and Yana Lishkova

Landing safely in unknown or changing terrain requires more than avoiding obstacles: a drone must approach a landing site without losing the ability to divert if that site becomes unsafe. This project validates an optimization-based filtering framework that modifies nominal trajectory commands only when required to preserve safety, steering, or contingency. Three-dimensional terrain maps are transformed using Poisson's equation into smooth safety fields, which are enforced through control barrier functions for hazard avoidance. Lyapunov-based certificates steer the drone toward a selected landing zone while defining certified regions from which alternative zones remain reachable. A combinatorial contingency condition maintains a prescribed number of these alternatives throughout the maneuver, enabling the landing target to switch when the active site becomes unsuitable. Model predictive control provides dynamically feasible nominal trajectories, while the certificate filters constrain their execution. The project evaluates alternative ways of integrating these components to establish a computationally efficient architecture for safe landing, diversion, and contingency preservation in Mars-analog environments

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.

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.

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. The extreme sensitivity of qubits may enable detection of vanishingly small changes in the local environment, a property applicable to molecular sensing. Towards this goal, molecular platforms, including open shell transition metal complexes (spin = ½), have garnered interest for their potential to produce measurable changes in spin relaxation behavior when exposed to environmental stimuli. Previous efforts in the Agapie Group have demonstrated how perturbations at or near the spin center are transduced into detectable changes in the electronic structure of a molecular qubit. Recent work expands this platform in two directions. Through the formation of crown-ether-functionalized copper qubits, the effects of cation binding on spin lattice relaxation has been explored. In parallel, a new sensing modality has been investigated using novel vanadium(IV) complexes supported by a *bis*-triazinyl bipyridine scaffold. Ongoing efforts target installation of variable ligands in the equatorial plane of these vanadyl qubit candidates, expanding understanding of the sensitivity of transition metal complexes.

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.

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.

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 .

Extensions of the Voronoi Implicit Interface Method to unstructured meshes with applications to grain boundary evolution within a 3D X-ray diffraction inversion algorithm
Eitan E. Camacho
Mentors: Kaushik Bhattacharya and Carter Cocke

The Voronoi Implicit Interface Method is an efficient algorithm for tracking multiple interacting phases whose motion is governed by complex physics. The method tracks any number of phase boundaries with a single unsigned distance function and uses a narrow band method to limit updates to a reduced domain near interfaces. Distances are calculated using the Fast Marching Method, an efficient algorithm for solving the Eikonal equation. While the initial Voronoi Implicit Interface Method was implemented on a Cartesian mesh, we extend the algorithm to unstructured tetrahedral meshes in the case of advection due to an externally given, time-independent velocity field. This is important since unstructured meshes provide more flexibility when working with complex geometries. The method is 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. The computational efficiency and accuracy of the method within the application are compared to a previously implemented naive multiple level set approach.

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 for space solar power applications
Enya Cao
Mentors: Harry A. Atwater, Jr., Andy Nyholm, Elsie P. Loukiantchenko, 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 ⟨110⟩ orientations generally produces smoother surfaces than along the ⟨211⟩ orientations, whileSi(100) results primarily reveal smooth, corrugation-free regions along the ⟨100⟩ orientations. Electroplating conditions remain ongoing work and are expected to reveal how stressor residual stress gives rise to distinct morphological features.

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.

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.

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.

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.

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.

Elucidating the design space of consistency flow maps for time-dependent PDEs
Lufang A. Chiang
Mentors: Anima Anandkumar and Jiachen Yao

Learning the evolution operator of time-dependent partial differential equations (PDEs) offers a direct route to fast and accurate surrogates for scientific simulation. Instead of composing many small solver or rollout steps, one can learn a transition kernel that advances the solution over an arbitrary time span in a single step. Recent methods realize this kernel as a consistency flow map along physical time, and train it by enforcing trajectory consistency equations. However, it remains unclear how these learning equations affect the model accuracy and long-term stability, and fair cross-method comparison is still limited. In this paper, we elucidate the design space of consistency flow maps, spanned by the form of the consistency equation and the source of the supervising velocity, and run a systematic study across PDE families. From this space we propose IconFlow, an integral consistency flow map trained with a decoupled, derivative-free multi-point objective. On representative PDE benchmarks, IconFlow leads the Pareto front in few-step prediction and stays robust under temporal extrapolation, providing a more efficient and reliable surrogate for time-dependent PDEs.

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 Carle 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.

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.

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.

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.

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 CO_2 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 CO_2 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.

Hardware-native quantum machine learning methods 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 exhibits stochastic brightness variations, unusual deviations from this variability are potentially indicative of rare or physically meaningful changes in the accretion system. Distinguishing expected damped random walk behavior from flares, repeating outbursts, or sustained state changes has broad implications for understanding the physics of black-hole accretion. This challenge is compounded by the millions of AGN light curves produced by modern time-domain surveys. This project investigates potential advantages quantum machine-learning methods provide to this problem, and circumscribes the cases and conditions under which they can compete with classical approaches. We adopted quantum variational rewinding with lagged-innovation and found comparable performance across anomaly strengths to one-class SVM and isolation forests. To develop an approach native to modern quantum hardware, we developed a novel quantum reservoir that maps two-band AGN light curves onto 128 data-encoding qubits and 28 additional interaction qubits using the native connectivity of IBM Fez. Across repeated noisy runs, the reservoir preserved 100% anomaly-ranking consistency, with quantum/classical score correlations in both bands. Rather than reproducing a classical classifier, this framework investigates quantum advantage through how quantum hardware represents and distinguishes light curves.

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.

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.

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.

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.

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.

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.

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. Crystals were polished to 50-100 μm and mounted on piezoelectric strain cells. Longitudinal and transverse voltages were measured from 1.7 to 10 K with a magnetic field applied along the c axis and uniaxial strain of up to ±0.6% (nominal) within the ab plane. In the first sample, the Hall signal varied most strongly at 1.7 K, changing linearly with strain; a linear extrapolation predicted a sign reversal of the Hall signal near -1.3% strain. Subsequent samples exhibited comparable linear tuning only at 5 K, while the response at other temperatures was substantially weaker. The next step is to establish data reproducibility and clarify temperature dependence before extending the study to larger strain amplitudes and different strain orientations within the ab plane.

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.

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 $\sim 1-2\times10^{39}\,\mathrm{erg\,s^{-1}}$. 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 $\log_{10}(\mathrm{age/yr})\sim 8.5$ 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 $0.15\pm0.09\,\mathrm{yr^{-1}}$.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.

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.

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.

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.

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

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.

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.

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.

Single-shot universal quantum computation with 4D lifted product codes
Ian C. Gill
Mentors: John P. Preskill, Madelyn G. Cain, and Muzhou Ma

Quantum low-density parity check (qLDPC) codes provide a low spatial overhead path to fault-tolerant quantum computation by taking advantage of nonlocal connectivity to produce high encoding rates. However, present schemes for universal quantum computation at finite size using high-rate qLDPC codes require a time overhead which scales with the code distance in order to protect against time-like measurement errors. In this work, we present a framework for performing universal quantum computation with low time overhead using univariate quadracycle (UQ) codes, which are a subclass of 4D lifted product codes. We identify high rate, high distance UQ codes with check weight 12, including examples with parameters [[306,66,15]], [[306,54,18]], and [[510,54,28]]. For these codes, we construct surgery gadgets which use the underlying metacheck structure of the UQ codes to require a number of logical cycles per Pauli measurement which is significantly below the code distance. Furthermore, we develop a single-shot CCZ state preparation and injection procedures for the high rate codes presented using lower rate UQ codes equipped with transversal logical CCZ gates, enabling universal quantum computation in conjunction with the native fold transversal gates and surgery gadgets. Together, these results establish UQ codes as a concrete finite-size route to universal fault-tolerant quantum computation that retains the low spatial overhead of high-rate qLDPC codes while substantially reducing the temporal overhead of their essential logical operations.

Why do LLMs misdescribe themselves?
Zakiriya H. Gladney
Mentor: Michael M. 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.

Investigating the characteristics of puncta formed by PPARγ
Anaya S. Gokarn
Mentors: Shasha Chong and Donghun Lee

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

Developing a new optical spectrometer to investigate the timing of ionization transitions relative to other plasma phenomena in the Caltech jet experiment
Miguel G. Guerra Zeferino
Mentors: Paul M. Bellan and Joshua Q. Morgan

This project aims to finalize the development of a new, fast optical spectrometer and then use this to measure the time evolution of the emission spectrum from the plasma in the Bellan Plasma Group's jet experiment at Caltech. In the Caltech jet experiment, a jet of plasma is created which develops a series of instabilities and eventually breaks apart. When the jet breaks apart, a ubiquitous plasma phenomenon called magnetic reconnection occurs, where the magnetic field rapidly reorganizes and releases energy into the plasma. X-rays are observed at this time, which is unexpected since it requires electrons to reach greater energies than were initially believed to be possible. Using data obtained from the new spectrometer, the timing of key transitions in the emission spectrum in relation to the emission of X-rays will be investigated to provide a high-temporal-resolution diagnostic of the plasma around the time of magnetic reconnection, which may be used by the lead plasma physicists to develop models explaining the X-ray bursts.

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 Ct is a smooth projective curve with an affine equation of the form Ct: $y7= x(x-1)(x-t)$ 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 Ct 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 $U_q(sl_2)$. 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 Solovyev 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.

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.

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 Gvadzabia
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.

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.

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.

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 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). In the Southern Ocean, water-masses are lightened and densified through water-mass transformation (WMT) (Abernathey et al. 2016), a key process in setting the strength of the meridional component of this circulation. Sea ice growth and melt dominate WMT rates in the Southern Ocean (Pellichero et al. 2018, Abernathey et al. 2016). However, little is known about the contribution of snow on sea ice to WMT rates due to the complexity of snow-sea ice-ocean interactions and scarce observations (Webster et al. 2018). While ECCO provides a combined freshwater flux including both sea ice and snow atop the ice, we used snow depth to individually estimate the snow freshwater flux. This revealed inconsistencies with the sea ice freshwater flux and observations that improved our understanding of the model's snow-on-sea-ice representation. With that knowledge, we assess how the snow freshwater contribution to WMT evolves with changing sea ice extent, improving understanding of the meridional component of the Southern Ocean overturning circulation.

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.

Verifying neural operators
Ryan L. Hsiang
Mentors: Anima Anandkumar and Robert Joseph George

Fourier neural operators (FNOs) learn mappings between functions but are evaluated numerically on finite grids, creating a gap between conventional grid-based verification and guarantees over continuous function spaces. We develop a sound interval bound propagation framework for certifying FNOs in the continuum limit. Our method applies sign-split interval propagation at a finite quadrature resolution and accounts for discretization error using analytically bounded Fourier kernels. To reduce the looseness of global estimates, we maintain cellwise enclosures of function values and spatial derivatives and use them to construct local quadrature corrections. These enclosures are propagated through spectral convolutions, pointwise lifting and projection layers, residual connections, and nonlinear activations. Bounds at grid nodes are then extended to the entire continuous domain using certified local Lipschitz constants, producing upper and lower-bound functions for every admissible input satisfying prescribed pointwise, supremum, and Lipschitz constraints. Numerical experiments on the FNO trained on the Poisson equation demonstrate the soundness and tightness of our approach. This framework provides a foundation for resolution-independent robustness verification of neural operators acting on continuous functions.

Probing atmospheric mass loss from WASP-33 b with metastable helium
Cheng-an Hsieh
Mentors: Heather A. Knutson, Morgan Saidel, and Pablo Drake

Gas giant planets on close-in orbits experience intense irradiation that can strip away their atmospheres. This atmospheric mass loss appears to be fastest for hot Jupiters orbiting early-type stars, but to date there are only two published observations of atmospheric escape from planets orbiting A stars. Here, we use the metastable He I 10830 Å triplet to search for evidence of an atmospheric outflow from the ultra-hot Jupiter WASP-33 b. We observed a single transit of WASP-33 b with Palomar/WIRC using a narrowband filter centered on this feature, which we jointly fit while accounting for stellar δ Scuti pulsations and observational systematics. We measure a helium excess depth of 3300 ± 1300 ppm, corresponding to 2.5σ evidence, and infer an outflow temperature of T_0 = 8500 +1248/−1638 K and $\log_{10}(\dot{M}/[\mathrm{g\ s^{-1}}])$ = 12.60 +0.44/−0.94. This mass-loss rate is in good agreement with the published value for WASP-33 b inferred from excess absorption in the Hα, Hβ, and Hγ lines (Yan et al. 2021). We jointly model the hydrogen and helium absorption using self-consistent mass-loss models and find that they underpredict the strength of metastable helium absorption relative to hydrogen. We conclude that helium may be a more useful tracer than expected for measuring atmospheric mass loss from planets orbiting early-type stars and highlight the need for improved outflow models.

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.

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.

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.

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

Electrolytic hydrogen offers a promising route to storing renewable energy, but current practice relies on platinum for the cathode reaction, whose cost and scarcity create a bottleneck to deployment at scale. Nickel-tungsten alloys are among the most promising non-precious replacements: nickel is abundant and already used in electrolysers, but binds hydrogen too strongly, and alloying with tungsten weakens that binding toward the Sabatier optimum. $Ni_{17}W_3$ is reported to reach 10 mA cm^{-2} at an overpotential of only 14 mV. Under operating conditions, however, the surface carries a substantial hydrogen coverage, and X-ray photoelectron spectra of the same material show tungsten in partially oxidised states alongside the metal. Here we use spin-polarised periodic DFT to build hydrogen coverage up site by site on pure Ni, on $Ni_{17}W_3$, and on $Ni_{17}W_3$ with two surface tungsten atoms each bearing an oxo oxygen, enumerating the symmetry-inequivalent sites at every step rather than assuming a filling order. We find that alloying and oxidation act principally by removing adsorption sites: of the sixteen face-centred-cubic hollows on pure Ni, ten remain accessible on the alloy, and oxidation displaces the additional hydrogen out of the hollow sites entirely. Because each surface reaches thermoneutrality at its own coverage, the descriptor that distinguishes these catalysts is where saturation falls. Density of states, charge analysis, and constant-potential calculations characterise the electronic origin of this behaviour, and nudged elastic band calculations probe the kinetics of hydrogen recombination.

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.

Strategic risk seeking in games
Oscar Jenner
Mentors: Eric V. Mazumdar, Boris Velasevic, and Nicolas Lanzetti

Strategic risk seeking, that is, optimism about other players' strategies, can be beneficial both in coordination games, where such optimism is natural, and in highly competitive games, where disadvantaged players must rely on opponent mistakes. In the spirit of recent work on risk aversion in games, this report studies strategic risk seeking in finite games and in linear quadratic dynamic games, focusing on when the proposed equilibria are unique and tractable to compute. In the first part of the report, a new risk-seeking equilibrium is introduced, combining risk seeking with bounded rationality. Independently of the game payoffs, conditions on the players' levels of optimism and bounded rationality are given that guarantee uniqueness of the equilibrium as well as convergence of best responses to it. In the second part, a risk-seeking equilibrium for linear quadratic dynamic games is proposed, for which existence and uniqueness guarantees are given. Numerical experiments show that mutual optimism can ease competition, but can also lead to free-riding. The open-loop formulation is then considered, where monotonicity conditions are given, paving the way for model predictive control applications.

Particle sedimentation in Newtonian and non-Newtonian fluids
Emily Jiang
Mentors: Sujit Datta, Mohamed Abdelgawad, and Craig Singiser

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.

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.

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.

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

Quantum key distribution (QKD) relies on the faithful transmission of entangled photons over long distances, where chromatic dispersion in optical fiber degrades temporal indistinguishability, reduces quantum interference, and limits secure key generation. This work investigates dispersion compensation for high-rate time-bin entangled photon sources as a means of preserving entanglement quality over long fiber links. Building upon previous classical-light characterization of a tunable dispersion compensator, coherent pulses were first used to calibrate optimal compensation settings before extending the study to pulsed and continuous-wave time-bin entangled photon sources. Entanglement visibility, coincidence rates, and Hong-Ou-Mandel (HOM) interference measurements were used to quantify the effectiveness of both local and nonlocal dispersion compensation over fiber lengths ranging from 50 to 200 km. The performance of the two compensation strategies was compared to determine their ability to restore temporal indistinguishability and preserve quantum interference after long-distance transmission. The results provide insight into the relationship between chromatic dispersion, photon indistinguishability, and QKD performance, while identifying practical operating conditions for high-rate entanglement distribution in fiber-based quantum communication systems. These findings contribute toward the development of scalable long-distance quantum networks and represent an important step toward realizing a functional quantum internet.

Constraining the critical Eddington ratio and variability signatures of changing-look AGN: A population-level study
Nora K. Kane
Mentor: Matthew J. Graham

Active galactic nuclei (AGN), powered by actively accreting supermassive black holes, are classified as Type 1 or Type 2 based on the presence of broad emission lines, traditionally attributed to viewing angle. Changing-look AGN (CLAGN), which transition between these types on unexpectedly short timescales at fixed orientation, challenge this picture, instead suggesting an accretion-state transition near a poorly constrained few-percent "critical" Eddington ratio whose precise value and link to CLAGN's variability remain unclear. Enabled by recent large-sky surveys, we convert Zwicky Transient Facility and Catalina Sky Survey photometry into archival-spectra-calibrated Eddington ratio curves for known CLAGN and a redshift/mass/luminosity-matched control AGN sample. Uniform Manifold Approximation and Projection subgrouping on variability and spectral features separates a stable, quiescent population from several more dynamic, variable subgroups, with CLAGN underrepresented among the quiescent group. The CLAGN population's Eddington ratio distribution constrains the critical threshold near 0.0275; most quiescent CLAGN sit at or below it, many for over a year. An empirical baseline built from control AGN further shows that variable CLAGN exceed typical AGN variability altogether, rather than merely sitting at its tail. Together, these results indicate changing-look behavior reflects a population-level shift in accretion state, motivating targeted spectroscopic follow-up of Eddington-ratio-selected candidates.

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.

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).

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.

Strain-engineering of MoS_2/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 MoS_2 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 MoS_2/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.

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 $M \sin i < 4.74$ MJ 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.

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

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 (Tc) 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.

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.

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.

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 fluid-structure interaction. 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 Ludwieg 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. Acoustic interactions within the dead-air region, including the generation of standing waves from the neck region and separation point are resolved and analyzed. Additionally, the interaction of the separation shock and shear layer are characterized as well.

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 and polarization anisotropies, inducing statistical mode coupling that can be used to reconstruct the lensing field. The resulting CMB lensing power spectra thus could serve as a robust probe for inference of cosmological parameters through a Gaussian likelihood which require an accurate modelling of the covariance matrix.

This project aims to compute the analytical bandpower covariance between lensing power spectra reconstructed with quadratic estimators from Planck, SPT-3G, and ACT-DR6, and to validate the calculation with Monte Carlo simulations. Neglecting cross-experiment covariance would produce unrealistically tight cosmological constraints, while overestimating it would weaken them. The autocovariances include both lensing cosmic variance and reconstruction noise, with additional mode coupling induced by survey masks. We focus on zero-mean Gaussian reconstruction noise uncorrelated with the lensing signal. Cross-covariance arises from overlapping sky coverage, where experiments observe the same lensing realization, and from correlated reconstruction noise. Preliminary results for the SPT-3G winter field show signal and noise cross-correlations of 23.4% and 12.5% with ACT, and 22.8% and 1.2% with Planck, respectively. We will extend this analysis to the SPT-3G summer and wide fields, where larger sky overlap is expected to increase these correlations.

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.

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.

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 mjlab. 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.

Synthesis and characterization of a trimetallic MNi_2 complex for CO_2 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 MNi_2 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 CO_2 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 $TiCl_3$ and $Ti(N\textit{t}Bu)Cl_2(py)_3$ 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 MNi_2 complex will be probed with single-crystal XRD, NMR, X-ray absorption spectroscopy, and computational studies. Finally, the aptitude of these complexes for CO_2 reduction will be studied.

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.

Isolating and quantifying the pre-targeting complex of the GET pathway in *Aspergillus fumigatus*
Avery C. Love
Mentors: William M. Clemons, Jr., 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 *Af*Get3 and *Af*Get4/5 to better understand the *A. fumigatus* pre-targeting complex and aim to quantify the effect of Retro compounds on the fungal GET pathway

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-Al_2O_3-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.

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 $LiHoF_4$
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 $LiHoF_4$ and probe quenched domain dynamics using a time-resolved ringdown technique. Hyperfine electronic-nuclear (^{165}Ho spin) coupling increases the critical field. Resonantly pumping the E_{32} Ho^{3+} 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 $LiHoF_4$ 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.

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.

Simulating and tracking domain evolution in the quantum Ising ferromagnet $LiHo_{0.65}Y_{0.35}F_4$
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 $LiHo_xY_{1-x}F_4$ 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 $LiHo_{0.65}Y_{0.35}F_4$ using a Metropolis Monte Carlo method with acceptance criteria adding a transverse-field-dependent tunneling term in parallel to the standard thermal term. I developed algorithms to identify domains and track their evolution by maximizing intersection-over-union between frames, formulated as a linear assignment problem. I found that the domain count is bimodal during magnetization reversal, that domain sizes are approximately power-law distributed, and that domain nucleation occurs earlier with a stronger transverse field.

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 SiO_2 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.

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) are explosive deaths of massive stars that have had their envelopes partially stripped. These 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. A search campaign by the Zwicky Transient Facility (ZTF) recently discovered a number of Type IIb SNe with very low nickel-56 mass, whose light curves are dominated by the shock-cooling emission. These transients are fainter than typical Type IIb SNe and could in some cases be undetectable with ZTF. As a consequence, the events may have been systematically overlooked. If so, the sampling of Type IIb SNe is incomplete, biasing population-level measurements such as nickel-56 mass, progenitor properties, and explosion mechanisms. Thus, this study aims to constrain the explosion physics of cooling-dominated Type IIb SNe and quantify the extent to which the Type IIb population extends to fainter and faster transients. From the observables of a sample of cooling-dominated Type IIb SNe candidates, we define the parameter space for a grid of radiation-hydrodynamic simulations. Using the SuperNova Explosion Code, the grid models these events to characterize the underlying population.

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 20–30 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.

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 melanogastor* 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.

Trimodal contrastive learning for joint conditional probability estimation in fluid dynamics
Huy N. Nguyen
Mentors: Andrew M. Stuart, Claire Valva, and Hojjat Kaveh

Standard multimodal contrastive learning frameworks, such as Contrastive Language–Image Pre-training (CLIP), rely on aligning pairwise embeddings across modalities. Applications include matching text-image pairs and Lagrangian data assimilation. When extended to three or more modalities, these methods involve optimizing a summation of pairwise alignment losses. However, pairwise objectives only model one-to-one conditional distributions $p(u \mid v)$ but fail to capture joint conditional distributions $p(u \mid v, z)$. This gap is especially significant in physical systems where individual modalities can carry marginal predictive power but provide strong information jointly.

To resolve this, we present a trimodal contrastive framework based on a novel trimodal similarity metric that directly quantifies the joint alignment among three latent representations. We then formulate a corresponding trimodal contrastive loss function and validate it in Gaussian test settings, where learned representations can be benchmarked directly against exact analytical conditional distributions. Finally, we apply our framework to fluid dynamics by recovering Eulerian velocity fields

given Lagrangian flow trajectories and passive tracer concentrations data. Our final goal is to accurately reconstruct the original streamfunction even with sparse or noisy data, with potential applications in data assimilation for fluid dynamics, extending to active tracers like temperature, salinity and other measurements.

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.

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 τSPECT 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. τSPECT 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. The effect of moving geometry on the phase space is also explored, with neutron cooling on the order of 1 neV observed.

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.

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.

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 tumorogenesis; 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.

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.

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 autonomous robot that combines aerial and ground locomotion using shrouded propellers surrounded by treads to traverse environments where terrain, access, and mission requirements are variable. However, its high acoustic signature may limit operation near people. This project investigates whether redesigning the structural spokes shrouds upstream of the propeller can reduce noise while preserving the geometry required for ground mobility. A single-propeller test rig was designed and fabricated to allow interchangeable shroud panels. Controlled design families vary spoke count while holding total blockage approximately constant, and include systematic comparisons of tapered, rounded, and various spiral spoke geometries intended to spread out and reduce interactions. Acoustic measurements are recorded with a fixed measurement microphone and sound pressure level meter at consistent motor operating points, and frequency spectra compare overall sound level, blade-passing tones, their harmonics, and prominent geometry-dependent tones. The resulting comparisons will provide a better understanding of the noise sources on M4 Phoenix and from upstream propeller shrouding in general, identify quieter spoke configurations, and guide a quieter shroud design for integration into the full M4 Phoenix robot.

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.

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.

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.

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.

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 (G) and the freezing-front velocity (V) dictate the finished materials properties. This project has focused on making a movable freeze-casting setup which allows the manufacture of much longer samples than previous static freeze-casting setups. The movable setup physically decouples G and V enabling the exploration of freeze-cast materials, crystal morphologies, and properties and regimes that are inaccessible to other freeze-casting apparatuses. After getting the machine operational the samples from the movable machine were validated to be identical in structure to those from the static freeze-casting machine at the same parameters. The machine then allowed exploration of different parts of the G, V graph, exploring the crystal morphologies and material properties at new G-V regimes. It also has allowed the exploration of much larger samples than the static freeze-casting machine can produce, allowing investigation for the potential scalability of freeze-casting in the future. Additionally, the physical decoupling has allowed the machine to form continuously frozen samples while altering the morphology with length.

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 λ = 0.75-5 μm 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.

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.

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.

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.

Establishing genetic methods for chemoreceptor manipulation in rove beetles
Conner I. Sakae
Mentors: Joseph Parker and Hayley Smihula

Pasadena, California is home to *Sceptobius lativentris*, a rove beetle species and obligate symbiont of the ant *Liometopum occidentale*. *Sceptobius'* 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 *Sceptobius*. 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.

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 to the devices. In the current Yb:YVO_4 platform, light is coupled into focused-ion-beam-fabricated photonic cavities via total internal reflection off 45° angled cuts, achieving efficiencies of only ~30% 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 finite-difference eigenmode (FDE) and finite-difference time-domain (FDTD) simulations. For YVO_4, the optimized design achieves a simulated coupling efficiency of 98.4%, representing a substantial improvement over the existing approach. We further extend the design to $CaWO_4$, a promising alternative host crystal, and to 45°-cut YVO_4 chips, obtaining similarly high simulated efficiencies of 98.1% and 98.2%, respectively. Finally, we evaluate the robustness of the optimized designs to alignment and fabrication errors, demonstrating that they remain efficient under realistic experimental conditions. 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.

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.

Validating a computational flow solver against experimental data for multi-fidelity aerodynamic shape optimization
Aritz Schube Barriola
Mentors: Steven Brunton, Matthew M. Gherman, Ana Larrañaga, and Patricia Garcia Caspueñas

This project supports the development of a multi-fidelity active learning framework for aerodynamic shape optimization, which requires understanding the cost and accuracy trade-offs among available flow-modeling tools. The work began with a technical comparison of four modeling tools spanning a wide range of computational cost and physical detail, from fast approximate methods to a general-purpose flow solver capable of resolving complex, three-dimensional turbulent flow. Based on this comparison, the general-purpose solver was selected for deeper validation. A repeatable workflow was developed for generating computational meshes around airfoil shapes and configuring the solver 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 speeds, and the results were compared against published wind-tunnel measurements. The next phase of work shifts from validation to aerodynamic shape optimization, and will involve comparing candidate optimization approaches before selecting a method for optimizing an airfoil shape to meet specified lift and drag targets.

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.

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 quadcopters 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 quadcopters with obstacles. We perform ablations over practical design choices, and compare with other methods such as PPO. We also gradually study the effects of practical relaxations of the theory's assumptions, and conduct some theoretical investigations to obtain wider bounds on the learning rate. Finally, we hope to implement the theory on a Crazyflie quadcopter.

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.

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~μm], 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~μm. 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.

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.

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.

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.

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.

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 CO_2 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, Fe_3O_4 (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.

Designing a superconducting microwave resonator to magnetically couple to $^{171}Yb^{3+}:CaWO_4$
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 $^{171}Yb^{3+}:CaWO_4$ due to the combination of long intrinsic coherence and a strong microwave transition dipole moment. By coupling a bulk $^{171}Yb^{3+}:CaWO_4$ 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.

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.

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.

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.

Developing improved computational algorithms for the stable Andrews-Curtis conjecture
Elliot D. Steffen
Mentor: Sergei G. Gukov

This project investigates whether classical computer search and machine learning can make difficult problems in group theory more tractable. It focuses on the stable Andrews-Curtis conjecture, which asks whether symbolic descriptions of the trivial group can always be simplified by elementary moves. The project develops a stable move that temporarily introduces an additional symbol and performs a change of variables. The resulting algorithm also separates ordinary moves from this more expensive transformation, using it only when it is likely to improve the search. On a 60-case benchmark, the best setting solved every case, reducing the average number of search steps from 26,130 to 9,754 and the average running time from 552.71 seconds to 48.18 seconds. It did not solve additional difficult cases and usually produced longer solution sequences. These results show that selectively scheduling stronger transformations can reduce computational effort on difficult instances that were already solvable. Future work will explore machine learning integration to improve search decisions and solution finding.

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 Hα 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.

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.

The Schauder fixed point property for countable amenable groups
Lizzie V. Teryoshin
Mentor: Omer Tamuz

An amenable group G is said to have the Schauder fixed point property if every action of G on a compact convex set has a fixed point. For example, the infinite cyclic group has the Schauder fixed point property while the infinite dihedral group does not. Work of Oliver from 1975 classified which finite groups have the Schauder fixed point property, but the situation for infinite groups is still not known in general. In particular, it is unknown whether the lamplighter group $\mathbb{Z}_2 \wr \mathbb{Z}$ has the Schauder fixed point property. In joint work with Nicolas Monod, we show that $\mathbb{Z} \times \mathbb{Z}_p$ (a quotient of the lamplighter group with lamps in $\mathbb{Z}_p$) and certain semidirect products by $\mathbb{Z}$, have the Schauder fixed point property.

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 coarse 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.

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 15N isotope labeling.

Surface plasmon resonance spectroscopy for investigating the growth of molecular additive layers on copper for CO_2 reduction
Cayley Y. Tolbert-Schwartz
Mentors: Harry A. Atwater, Jr., and Ellis Spickermann

The electrochemical CO_2 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 CO_2 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 CO_2 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.

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.

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.

Breaking orbital degeneracies of directly imaged companions using KPIC relative radial velocities
Tyler N. Tuttle
Mentors: Dimitri P. Mawet and Clarissa R. Do O

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 ($a \approx 464$ AU, $P \approx 5321$ 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.

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.

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 a task where the robot must track a moving target with the head while an MPC-CBF safety filter governs body navigation around obstacles. We have developed two parallel control stacks sharing the same perception and head hardware: a model-based stack using visual servoing and an MPC-CBF planner, and a learning-based stack in which a single distilled recurrent policy outputs both gaze and body commands, integrating perception directly into training. Both stacks have been tested in simulation with the head assembled on the robot across several detectors (YOLOv11n, YOLOE-26). In simulation, the pan-tilt head reorients toward the target faster than the body can turn, keeping the target in view without waiting for a full-body rotation. Early results also suggest the learning-based stack is currently sensitive to detector latency; improving this through training, along with a full comparison across detector configurations and hardware trials with the real head, are ongoing.

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 *Sceptobius lativentris*, which is less genetically tractible 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.

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.

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.

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.

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.

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.

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.

Developing a molecular replacement method for phasing low-resolution MicroED data of paramagnetic transition metal complexes
Kyler E. Williams
Mentors: Hosea Nelson and Sam 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.

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.

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.

Streaming state-space hybrid mesh recovery model
Xinran Xie
Mentors: Pietro Perona and Ilona Demler

We present Streaming State-Space Hybrid Mesh Recovery Model, an efficient streaming HMR with high reconstruction quality. Traditional video HMR is accurate but offline: the temporal module attends across the whole clip, which makes the clip a prerequisite for any single output frame and makes memory grow with the history kept. Hybrid state-space architectures remove such cost in language models, but a language model is already causal, so substituting a recurrent layer preserves what every token had. An offline mesh recovery model is bidirectional, and the same substitution removes the future along with the quadratic term. We resolve this causality mismatch by a retraining-free state-space replacement pipeline which enables streamability without significant loss in reconstruction quality, using only 8.5k frames from three benchmark datasets for fine-tuning. We achieve **(a). higher runtime and memory efficiency**. The converted model carries 1.87 MB of state per tracked person with no key-value cache, unchanged from 1,000 to 131,072 streamed frames and identical at every context window measured, against 25.2 times more state for an attention reference at a 960-frame context. Full replacement raises the processable clip length to 127k frames, by 8.3× the offline backbone's. We also achieve 1.39x and 1.15x speed up against the two streaming baseline models Human3R and Online HMR with higher performance. **(b). preserved accuracy on both in-domain and zero-shot tasks**. With fine-tuning, at 50% replacement, the hybrid beats the offline backbone it is derived from, and it beats SOTA streaming model baselines on all three evaluation sets. Furthermore, we propose not only a single model, but also a transferrable recipe for making a bidirectional mesh recovery architecture efficient, streamable, and high-quality.

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 Heinaevaara

The tracial joint spectral measure is a measure over R^2 and originally introduced by Heinävaara to prove Hanner's inequality. We are interested in the measure itself: its continuous part, μ_c. μ_c is defined using two Hermitian matrices A and B, and associated with eigenvalues, although not the eigenvalues of A and B. We already know μ_c is compactly supported beforehand, and when A and B commute, it is zero everywhere. We perturb this trivial case a little bit. We let A and B start from diagonal matrices and apply an $O(\varepsilon)$ perturbation; then study the location, shape, and density of μ_c. For the location, we proved μ_c is supported near several collision lines, where two eigenvalues are equal. Near the collision lines, μ_c can be approximated by submatrices of A and B. The shape and error of approximation are tied to the size of the submatrices. 2×2 approximation is the simplest yet very crude: geometrically it gives a union of several ellipses and the error is at $O(\varepsilon)$. Larger sizes give reduced error but also have more complicated shape. Our questions on the shape and density of μ_c are answered by the submatrix approximation.

Optimizing sequential resin–latex extractions from guayule through analytical characterizations

Ruoyun Xu
Mentors: Julia A. Kornfield and Florence Mueller

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.

Integration of PL spectroscopy and twisted light-assisted STM for nanoscale exciton research on ML-MoS_2

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 MoS_2 (ML-MoS_2). Gold-assisted transfer is investigated for preparing large-area ML-MoS_2. 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 MoS_2 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-MoS_2.

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.

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 Fe_3Sn_2, 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.

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.

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.

Transfer learning for AI-assisted mathematical discovery in combinatorial group theory
Caroline Zhang
Mentors: Anima Anandkumar and Robert Joseph George

AI is transforming mathematical research, with LLM and RL-based solvers finding proofs to previously unsolved problems. Motivated by these recent advances, and by how human mathematicians are able to transfer knowledge between problems, we investigated and enabled AI systems to learn general mathematical representations that can be applied across different problems. We focused our study on 24 families of problems in combinatorial group theory, including the notable Andrews–Curtis conjecture, an open problem for over 60 years. We developed a reinforcement learning (RL) framework and transfer learning pipeline, and we further trained a transformer-based foundation model capable of pretraining on all 24 families in the benchmark. Our results show that (1) transfer learning is effective across diverse group theory problems, yielding 17 robust source-to-target pairs with positive transfer after 1M training timesteps; (2) successful transfer can be attributed to both reusable learned representations and fine-tuning during training; and (3) pretraining on different source problems enables agents to solve complementary subsets of benchmark instances, resulting in greater dataset coverage than training from scratch alone. Together, these findings suggest that shared representations can emerge across problems in group theory and that training on one task improves performance on others, highlighting the potential of generalizable RL agents for mathematical solving and discovery.

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 5.0×107-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 and 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 infleuences the establishment and maintenance of beneficial host–microbe interactions.

Accretion flow and jet evolution during the 2021 outburst of GX 339−4
Yuchi Zhang
Mentors: Fiona A. Harrison and Shina Adegoke

GX 339−4 is a stellar-mass black hole system that brightens in recurrent outbursts as gas from the companion star changes the structure of the surrounding accretion flow. This project investigates how the disk, the coronal high-energy emission, and the radio jet evolved during its 2021 outburst. To do this, we used data from the Neutron Star Interior Composition Explorer mission, covering the 0.3—12 keV energy range, together with contemporaneous radio measurements from MeerKAT between January and November 2021. Preliminary results show that while in the disk-dominated soft state, the disk-luminosity versus temperature relation in the GX 339-4 system agrees with theoretical prediction, suggesting that the disk likely remained stable at the innermost stable circular orbit of the black hole throughout the soft state evolution.

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.

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.