

Alkyne synthesis by the cross-coupling of nitroalkanes with nitroalkenes

Adais M. Arora

Mentors: Marvin Parasram and Brian M. Stoltz

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

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

Michelle Cai

Mentors: Mikhail G. Shapiro and George H. Daglian

Abstract withheld from publication at mentor's request.

On-policy self-distillation for physically plausible video generation

Hongqiao Chen

Mentors: Georgia Gkioxari and Ziqi Ma

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

Heritable targeted DNA insertion using the R2Tg retroelement complex

Lucas N. Garcia

Mentors: Gözde S. Demirer and Kimberley T. Muchenje

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

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

Ido Geffen

Mentors: Ravinder S. Dahiya and R. Michael Alvarez

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

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

Designing MOF based catalysts through electron density descriptors

Shrish Goel

Mentors: Anima Anandkumar and Danish Khan

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

Virtual Merge models in Hopf algebra Markov chains

Nikash Gupta

Mentor: Matilde Marcolli

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

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

William Jouayed

Mentors: Emma Lessieur and Melany L. Hunt

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

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

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

Elianna G. Kondylis

Mentors: Gregory Stephanopoulos and Rustem F. Ismagilov

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

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

Henry B. Lane

Mentors: Paul D. Asimow and Carl Swindle

3D magnetotelluric (3DMT) forward modeling has consistently posed itself as a computationally demanding problem. Despite its wide range of applications, ranging from geothermal and mineral exploration to studies of crustal and mantle structure, the high computational cost of such simulations remain prohibitive for developing large-scale inversion methods and uncertainty quantification methods. Building on recent advances for solvers that combine high-order Nédélec-type finite elements with the auxiliary-space Maxwell (AMS) preconditioner, this work develops a complex-valued, GPU-native AMS solver and constructs a modified high-order 3DMT solver around it, eliminating the need for the real-valued block decomposition. Moreover, we validate the solver by comparing it against other validated solvers and accepted numerical benchmarks, demonstrating accurate MT responses. Finally, we demonstrate that the solver has improved convergence behavior relative to the reference solver baseline, when implemented on GPUs inherits the favorable scaling behaviors of the CPU-based reference solver, and show that the associated wall-time is significantly reduced relative to the CPU implementation of the reference solver. Together, these developments provide a high-accuracy, geometry (bathymetry/topography) conforming, and computationally convenient and scalable foundation for next-generation 3DMT inversion.

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

Alexis N. Lindenfelser

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

High entropy oxides (HEOs) are an emerging class of materials with applications in catalytic conversion and energy storage. By combining five or more metal cations in equimolar amounts, HEOs have high configurational entropy that stabilizes single-phase solid oxide solutions, enabling enhanced tunability and offering distinct active sites to facilitate multistep reactions, such as the reverse water gas shift (RWGS) reaction. Experimental studies on spinel-type MnFeCoNiCuO_x HEOs found that Co- and Mn-rich samples (type 1) had enhanced RWGS activity compared to Fe-, Cu-, and Ni-rich samples (type 2), despite similar configurational entropies. We hypothesize that local mixing effects from distinct cation coordination environments influence the reversibility of oxygen vacancy formation and thus modulate RWGS activity. In this work, we use partially inverse CoFe_2O_4 spinels, generated using special quasirandom structures, as a computationally tractable structural analog for MnFeCoNiCuO_x HEOs. Mn, Ni, and Cu are substituted into local coordination shells around surface oxygen sites in (100) slab models to simulate type 1 and type 2 atomic environments. We compute oxygen vacancy formation energies using density functional theory (DFT) calculations, demonstrating that near-neighbor effects explain the activity discrepancy between samples, and showing that HEOs can be tuned by elemental compositions to improve catalytic activity.

Continuous best-response learning for direct exploitability descent

Eric Y. Lu

Mentor: Tony Yue Yu

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

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

Xingchuan Ma

Mentors: Sarkis K. Mazmanian and Daniel Dumitrescu

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

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

Yingying Ma

Mentor: Mark Simons

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

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

Prince Mutegetsu

Mentors: Christian Frankenberg and Eva-Marie Metz

California forests are increasingly affected by disturbances such as wildfire, drought, and heat. Wildfires, in particular, can prevent forests from recovering after disturbance, leading to conversion from forest to shrubland. However, it remains difficult to determine which conditions drive these conversions and which forest areas are most at risk. This project adapts a data pipeline to examine whether forests show signs of slower recovery following disturbance across California. The pipeline uses monthly vegetation observations from the Landsat mission spanning 1984 to 2025, removes seasonal patterns and long-term trends, calculates temporal autocorrelation (TAC) and variability in rolling windows, and then estimates their temporal trends using Kendall's tau. These calculations are computationally expensive, so I used Numba to parallelize the pixel-wise computations across 48 CPU cores, enabling processing of 15.2 million pixels at 90 m resolution across the Sierra Nevada in 17 minutes. The resulting maps were compared with known forest-to-shrub conversion sites and Forest

Inventory and Analysis (FIA) field plots, where the observed patterns generally matched. These results demonstrate that the pipeline can be scaled from small test regions to a California-wide analysis. Future work will strengthen validation and support development of a model that identifies forests at risk before conversion.

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

Rachel E. Nazareth

Mentors: Dimitri P. Mawet and Andrea Lin

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

Structural analysis and reconstitution of RNA polymerase II ubiquitination

Jia-Xin Nien

Mentors: Isaac Fianu and Anjali Mirmira

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

Critical mineral recovery via electrosorption across zeolite films

Cleatus R. Oakes

Mentors: Pavel A. Kots and Mark E. Davis

LTA zeolite films are investigated as selective ion-exchange and molecular-sieving materials for sustainable lithium recovery from brines. Conventional solar-evaporation extraction is increasingly viewed as environmentally and operationally unsustainable, motivating the development of more efficient separation strategies. LTA zeolites, aluminosilicate crystals with a well-defined cubic framework and uniform pore sizes, can be synthesized as thin, continuous films on metal substrates and offer size- and hydration-based ion selectivity. In this work, LTA films were grown on nickel supports via hydrothermal crystallization. The resulting coatings were characterized for thickness, continuity, and adhesion. Preliminary ion-exchange experiments using Li⁺/Mg²⁺ control solutions were conducted to quantify intrinsic lithium uptake and evaluate competitive ion exclusion. The

zeolite-coated electrode was then integrated into an electrochemical cell containing a $\text{Li}^+/\text{Mg}^{2+}$ control solution. Baseline transport was measured under open-circuit conditions, followed by field-driven experiments. Preliminary results show that LTA zeolite exhibits a high decationization capacity enabling strong lithium selectivity, indicating its potential as an efficient, scalable membrane material for environmentally sustainable lithium extraction. Further optimization of film thickness and applied field strength may enhance selectivity and recovery, advancing the development of more efficient lithium extraction technologies.

Evaluating the accuracy and stability of automatic differentiation of eigenvalue computations

Maria T. Papazoglu

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

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

Wave attenuation and hydrodynamic analysis of artificial rocky reefs

Kyra T. Phaychanpheng

Mentors: Christine M. Baker and John O. Dabiri

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

Creating benchmarks for robust space robotics decision making with reinforcement learning

Ethan Pichon

Mentors: Benjamin Riviere and Gunter Niemeyer

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

space reproduces simulation exactly, while an otherwise identical policy trained on privileged state loses half its performance. These benchmarks establish baselines for an online RL planner that uses state-space mixing to explore bottleneck states.

Transition from inspiral to plunge for eccentric EMRIs with spinning secondary

Oleksandra Pyshna

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

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

Directional Gaussian smoothing for gradient descent

Ryan M. Rodrigue

Mentors: Federico Pasqualotto and R. Michael Alvarez

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

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

Brittany M. Rosario

Mentor: Aaron D. Ames

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

full-body pose from the combined inertial and ranging data and stream it into a humanoid simulation in place of the virtual reality input, giving the project a portable, low cost motion source and, ultimately, a testing ground for developing prosthesis controllers before they reach hardware.

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

Trisha A. Sabadra

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

Ground-based high-resolution spectroscopy resolves a planet's spectrum into individual atomic and molecular lines, letting us characterize the composition and dynamics of exoplanet atmospheres. High-resolution cross-correlation spectroscopy has recently begun to resolve the dayside atmospheres of ultra-hot Jupiters in detail, including WASP-189b, whose dayside emission spectrum Sanchez et al. recently characterized with IGRINS, detecting Fe I, Mg I, Si I, H₂O, CO, and OH. A dayside detection alone, however, cannot reveal how heat absorbed on a planet's permanent dayside is transported to and processed on its night side. The principal objective of this project is to determine whether the same species detected on WASP-189b's dayside are also present in its night-side and terminator atmosphere. To do this, we adapt PANCAKE, the high-resolution cross-correlation spectroscopy pipeline introduced by the KPF SURFS-UP Survey for the Keck Planet Finder, to IGRINS-2, the upgraded near-infrared spectrograph at Gemini North, and obtain night-side emission spectroscopy of WASP-189b (Program GN-2026A-Q-114). We first validate our approach against a known result, using the publicly available pipeline IGRINS_transit to successfully replicate Sanchez et al.'s published dayside Fe I detection of an archival IGRINS eclipse dataset of the same target (GS-2023A-Q-231). Applying our own adapted pipeline to our primary dataset, however, we do not yet recover a convergent signal at WASP-189b's expected orbital velocity for any of eight atomic and molecular species searched, and report a tentative non-detection for the planet's night side.

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

Efe Sakarya

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

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

Task-driven models for neural speech representations

Yourui Shao

Mentors: Francis R. Willett and Yisong Yue

Clear characterizations of how sequences and language are represented in the brain are crucial for developing targeted decoding approaches for speech brain-computer interfaces. Given neural recordings characterizing low-level phonemic and higher-level syntax representations in the brain, we

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

Reflector-solver co-training for large language models

Arjun Sharma

Mentors: Yisong Yue and Atharva Sehgal

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

Computational modeling of nitrogen reduction intermediates on LaFeO₃ and LuFeO₃ perovskite oxide surfaces

Daniella E. Sherman

Mentors: William A. Goddard III and Sejun Kim

This study used density functional theory calculations to investigate nitrogen adsorption and nitrogen reduction behavior on LaFeO₃ and LuFeO₃ perovskite oxide surfaces. The goal of this project was to determine how changing the A-site cation of these perovskite oxides from lanthanum to lutetium affected surface adsorption, nitrogen activation, reaction intermediates, and competition with the hydrogen evolution reaction. First, surface models were constructed and optimized. Nitrogen adsorption was then investigated at multiple surface sites. Nitrogen adsorption calculations were followed by calculations of hydrogenated nitrogen intermediates and adsorbed hydrogen structures. Binding energies were calculated relative to the clean surfaces and reference molecules to compare the varying stability of the different evaluated adsorption geometries. The results of this project showed that nitrogen adsorption depended strongly on both the adsorption site and surface composition. Particularly, LaFeO₃ and LuFeO₃ favored different adsorption geometries and showed differences in adsorption strength and intermediate stability across various calculations. These findings indicate that changing the A-site cation affects the surface environment of perovskite oxides and can provide a future strategy for tuning perovskite oxide catalysts for nitrogen reduction and sustainable ammonia production.

Sequential quantum circuits mapping between symmetry-protected gapless models

Shreyas Singh

Mentors: Xie Chen and Vibhu Ravindran

Matter at low temperature often exhibits interesting quantum effects by virtue of occupying a low-energy state of its Hamiltonian, which dictates energy interactions between adjacent particles. Consequently, a phase transition between two systems is detectable by continuously interpolating their Hamiltonians and searching for a discontinuous change in some parameter. An alternative (and largely equivalent) method for classifying phases is inspired from quantum computing, wherein two systems are in separate phases if they cannot be connected by a finite-depth circuit of local unitary operators, instead requiring a linear-depth "sequential" circuit. This latter scheme has successfully classified phases among one-dimensional "gapped" models, whose second-lowest energy state remains separated from the lowest as the system size grows. We investigate the potential of the sequential circuit formalism for classifying gapless models, which has thus far remained unclear. In particular, we construct an invertible sequential circuit mapping between the two gapless models $X+ZZ$ and $ZXZ+ZZ$, as well as a noninvertible circuit mapping between the critical Ising models $X+ZZ$ and $X+YY$, both protected by $\mathbb{Z}_2$ and time reversal symmetries. The existence of sequential circuits performing these mappings, but not of finite-depth circuits, supports the utility of sequential circuits for classifying gapless phases.

Ruling out an intermediate mass progenitor donor in V404 Cygni

Joe F. Smith

Mentors: Kareem J. El-Badry and Cheyanne Shariat

Typical black holes (BHs) in BH low-mass X-ray binaries (BH LMXBs) are understood to not accrete a significant amount of mass over their lifetimes. This assumption can now be tested through binary evolution accounts due to the recent discovery that the BH LMXB V404 Cygni belongs to a hierarchical triple system, which provides an age estimate of 3-5 Gyr via the tertiary. We employed `MESA` stellar evolution models for V404 Cygni across grids of initial donor masses and orbital periods and robustly examine influences on the system due to the CARB and RVJ magnetic braking prescriptions, non-conservative mass transfer thresholds via β , and Eddington limited versus uncapped retention. It is found that simulations with CARB generally allow higher mass progenitors than RVJ. Of the satisfactory runs, substantial BH growth occurs for moderately non-conservative mass transfer but models with $\beta \geq 0.8$ fail to reproduce V404 Cygni. However, while the simulations demonstrate that higher mass (2.5-3.0 $M_\odot$) progenitors can reproduce V404 Cygni's present state, they evolve to these states much earlier than is allowed by the age constraint provided by the tertiary. Therefore, our calculations favor low mass progenitors near 1.5-2.0 $M_\odot$.

Effects of IFN γ signaling on Treg phenotype, transcriptional state, and suppressive function

Amina Tajammal

Mentors: Andrew Luster, Ellen Rothenberg, and Alexis Yero Diaz

IFN γ is a dominant cytokine produced during influenza infection and signals through the IFN γ receptor (IFNGR1/IFNGR2 complex) to activate STAT1-dependent transcriptional programs. Tregs express IFNGR1 and are therefore directly responsive to IFN γ . Prior studies demonstrate that IFN γ signaling promotes T-bet and CXCR3 expression in Tregs, facilitating adaptation to type 1 inflammatory environments. However, sustained IFN γ exposure may reprogram Tregs away from a GATA3⁺ type 2-regulatory phenotype, potentially impairing their ability to control allergic inflammation. Dissecting Treg-intrinsic IFN γ signaling is thus essential to determine whether IFN γ directly drives Treg reprogramming and functional impairment during viral-induced asthma exacerbation. We propose that RVIs drive Treg reprogramming, thereby reducing their capacity to suppress type 2 inflammation. This project aims to understand how viral-induced cytokines, particularly IFN γ , reshape Treg phenotype and potentially immunosuppressive functions, which is critical for identifying therapeutic strategies to restore Treg function, prevent exacerbations, and improve outcomes for patients with asthma.

Mitochondrial inheritance and fertilization-driven asymmetry in the early mouse embryo

Katherine W. Tam

Mentors: Magdalena D. Zernicka-Goetz and Breanna M. McMahon

With global rates of infertility increasing, the source of the earliest patterning and developmental potential of the mammalian embryo continues to be questioned. Despite longstanding beliefs of development starting with largely equivalent blastomeres, growing evidence suggests that asymmetries emerge as early as fertilization, but how these initial differences are generated and propagated through cleavage and ZGA remains unclear. Here, we investigate whether mitochondrial organization, maternal transcriptional organization, and cytoplasmic dynamics contribute to early symmetry breaking in the mouse embryo. With the use of fixed and live imaging of oocytes, zygotes, and cleavage-stage embryos, we are able to track maternal and paternal mitochondrial populations and the distribution of transcripts in response to fertilization. Preliminary live imaging reveals that mitochondria becomes spatially asymmetric in the zygote, thus differentially partitioning between sister blastomeres which compound with future divisions. We further examine how these differences relate to fertilization-induced cytoplasmic flows and the asymmetric distribution of maternal factors, including YAP1 and transcripts associated with the zygotic genome activation. Together, these findings suggest that fertilization initiates intracellular heterogeneities that may provide an early bias from which developmental asymmetry is progressively amplified.

Self-driving imaging system for fetal brain magnetic resonance imaging (MRI)

Idil A. Turasi

Mentors: Polina Golland and Babak Hassibi

Abnormal fetal brain development must be detected and accurately characterized during the second or third trimesters of pregnancy for effective prenatal intervention. Because standard ultrasound examinations do not provide sufficient contrast and resolution for identifying smaller abnormalities and accurately measuring malformation size and severity, radiological assessment stands to greatly benefit from high-resolution T2-weighted volumetric MRI of the fetal brain. However, uncontrollable fetal motion during the scan precludes the diagnostic utility of the acquired images, necessitating repeated re-acquisition of motion-corrupted images. This project aims to improve a system that tracks fetal head motion and adjusts the next image slice in real time. I modified the existing segmentation model to identify spin-history artifacts separately from the brain, eyes, and background. I also constrained simulated artifacts to the brain region and tested different training weights. The best model achieved an average brain Dice score of 0.865 compared with 0.872 for the original model without artifact segmentation. Future work will use Gaussian regression of the distance from each voxel to the artifact plane and uniform normal sampling to improve artifact recovery.

Learning-based methods for ground-effect estimation on thrust-vector-controlled rocket vehicles

William A. Wrolstad

Mentors: Mory Gharib and Jack N. Caldwell

Precise vertical landing of small rocket vehicles requires accurate control in the presence of aerodynamic disturbances. As these vehicles approach the ground, aerodynamic ground-effect forces become a significant source of disturbance, and accounting for these forces in the controller may improve landing performance. This project advances the Daedalus family of electric rocket simulators as a test platform for autonomous precision landing algorithms. During the project, electrical systems for the next-generation Daedalus3 vehicle were developed, and the baseline waypoint controller for the Daedalus2 platform was implemented to support autonomous flight testing. This project also implements learning-based control methods to estimate and account for ground-effect forces on Daedalus2. Previously, similar learning-based methods have been implemented on quadcopters, where the thrust vector is fixed relative to the vehicle orientation, enabling ground-effect force estimation from the vehicle state alone. In contrast, thrust-vector-controlled vehicles decouple thrust direction from vehicle orientation, requiring additional information, such as the thrust vector control

outputs, to accurately predict ground-effect forces. The resulting model will be integrated into the landing controller on Daedalus2 and evaluated for its ability to improve touchdown precision. This work extends learning-based ground-effect estimation to a broader class of vertical-landing vehicles and provides a foundation for future testing on the Daedalus3 platform.

What if we trained video models on Mars?

Charlie Q. Yang

Mentors: Pietro Perona, Raphaela H. Kang, and Sevan Brodjian

Recent work suggests that foundational video models can develop an intuitive understanding of physical laws. Building upon this, we investigate the capacity of video models to internalize fundamental physical laws by training them in simulated environments with altered physics. To establish a baseline for evaluation, we introduce methods to quantify a model's intuitive physics understanding, including linear probing of Transformer layers to detect impossible physical trajectories and a representation prediction approach (VJEPA-H). We generated synthetic datasets representing universes with distinct physical properties, such as inverted gravity, zero friction, and collisionless environments, to train and evaluate models like V-JEPA and VideoMAE.