Student-Faculty Programs

Session L Abstracts

Real-time occupancy map generation from onboard LiDAR data for safety-critical control
Raffi A. Khondaker
Mentors: Aaron D. Ames, Gilbert Bahati, and Ryan M. Bena

Safe navigation in dynamic, cluttered environments demands onboard perception that maintains a timely, reliable model of nearby obstacles. LiDAR (Light Detection and Ranging) is widely used for its accurate range geometry and robustness to lighting, yet it is susceptible to occlusion, measurement noise, and ambiguity from ego-motion. This work presents a real-time occupancy-grid mapping approach from onboard LiDAR for safety-critical control. The method maintains a probabilistic two-dimensional grid and updates it with each scan in real-time. Free space is labeled along each viewing direction up to measured surfaces, which are marked as occupied, while the prior grid is propagated using uncertain robot motion estimates. Mild temporal decay and spatial smoothing limit overconfidence and reduce isolated false positives. The resulting grid is thresholded to obtain obstacle regions and also defines the domain for a Poisson-based Control Barrier Function, which computes a smooth safety field whose zero level marks obstacles. This mapping integrates into a safe navigation pipeline to adjust control inputs in real time on resource-constrained robots and motivates continued work on handling moving obstacles and adaptive thresholds.

LIP-guided reinforcement learning for robust and efficient bipedal locomotion
Timothy Kennedy
Mentors: Aaron D. Ames and Kejun Li

State-of-the-art locomotion controllers for bipedal robots often fall into one of two categories: model-free reinforcement learning (RL) controllers and model-based feedback controllers. Both of these options come with downsides, where RL controllers can be computationally expensive and have a high sim-to-real gap, whereas model-based controllers often lack robustness to unmodelled dynamics and terrain variations. We propose a hybrid framework that embeds a reduced-order linear inverted pendulum (LIP) model into the RL reward through a control-Lyapunov-function (CLF) term. The CLF steers learning towards LIP-consistent trajectories while preserving RL's flexibility. On the Unitree G1 humanoid, our method reduces cost of transport by 34% and position error by 75% relative to pure RL, outperforming both baselines in robustness, accuracy, and energy efficiency.

Integrating and characterizing the Unitree D1 robotic arm for whole-body control research
Eloise Zeng
Mentors: Aaron D. Ames and Zachary Olkin

This project focuses on integrating the Unitree D1 robotic arm into the Caltech AMBER Lab's robotics research interface, Obelisk, with the goals of enabling whole-body control experiments on the Unitree Go2 quadruped-armed system and improving the alignment between simulated and real-world robot behavior. After extensive debugging of the Unitree D1 SDK and Obelisk, I developed a teleoperation system consisting of a controller, state estimator, and hardware interface. The controller leverages the Pinocchio Python library to perform inverse kinematics on the gripper, and I implemented a joystick-based interface for intuitive user control. Despite these efforts, the D1 arm consistently failed to follow commands reliably, even after extensive troubleshooting and communication with Unitree support, ultimately leading to the conclusion that the hardware is defective. The project has since pivoted toward preparing for future research in model predictive control and reinforcement learning.

Benchmarking control barrier function methods for real-time resilient UAV swarm coordination
Edward S. Ju
Mentors: Aaron D. Ames, Pio Ong, and Ryan M. Bena

Swarm networks of unmanned aerial vehicles (UAVs) operating in formation often require communication connectivity while avoiding collisions. We present a study on real-time connectivity maintenance using control barrier functions (CBFs) and demonstrate implementation on Crazyflie hardware. We compare three CBF approaches: CBF-QP with naive affine connectivity constraints, CBF-SDP with linear matrix inequality Laplacian constraints, and CBF-QP with local connectivity conditions. Through simulation and hardware experiments, we analyze the trade-offs between computational complexity and task performance for each method. Our scalability tests demonstrate that our CBF-SDP approach achieves safe, real-time connectivity preservation. We provide quantitative task performance metrics comparing our methods with existing approaches, demonstrating real-time feasibility and scalability for swarm applications.

System identification to improve sim-to-real transfer of reinforcement learning policies for humanoid locomotion
Logan A. Woudstra
Mentors: Aaron D. Ames, Blake Werner, and Lizhi Yang

Reinforcement learning (RL) policies for humanoid locomotion are trained in simulated environments and typically experience performance degradation when transferred to hardware. This decreased performance is due to differences that exist between the simulator and reality, called the sim-to-real gap. One approach to minimize this gap is system identification, which seeks to estimate the physical parameters of the real-world system so that the simulator can better approximate its dynamics. In this work, we perform system identification, utilizing Wasserstein Distance and Maximum Mean Discrepancy to measure the distributional differences between trajectories rolled out in Nvidia IsaacSim and on hardware for the Unitree G1 humanoid robot. To search for the parameters that optimize these similarity measures, we employ the Covariance Matrix Adaptation Evolutionary Strategy, enhanced with parallelized generation evaluations for increased efficiency. With the development of our end-to-end system identification pipeline, we perform both sim-to-sim and sim-to-real alignment to demonstrate the effectiveness of our approach to learn parameters that enhance the transferability of RL policies between training and deployment environments.

AprilTag-based vision system for arm-guided object probing with a RealSense depth camera
Randolph L. Douge
Mentors: Joel W. Burdick and Yacine Derder

In order for robots like the Multi-Modal Mobility Morphobot (M4) to handle real-world fieldwork tasks, such as probing high-voltage electrical boxes or diagnosing solar array faults, they need a reliable method of sensing and interacting with objects in their environment. One critical missing piece in M4's design is a fine-manipulation system that allows for accurate, autonomous interaction with physical components. My project directly supports this capability by building a vision-based spatial recognition system using an Intel RealSense D405 depth camera, paired with AprilTag fiducial markers to serve as visual reference points in 3D space. Over the course of my work, I developed a Python-based pipeline that allows the RealSense camera to detect AprilTags in real time, extract their XYZ coordinates using calibrated depth data, and link this information to motion commands for a custom-built robotic arm. Unlike many robotic systems that rely on inverse kinematics or preset poses, this setup simplifies the movement logic by directly mapping camera-derived coordinates to arm positioning actions. The probing action is defined not just as reaching a point but also as confirming the presence of the tagged object, which reflects the kinds of spatial reasoning M4 will need in deployment. While developing this system, several challenges had to be overcome, including tuning the depth camera to deliver reliable distance measurements and ensuring smooth integration with OpenCV-based tracking. To test the system's robustness, a modular rail-based test stand was built to simulate a variety of object positions and tag orientations. The resulting pipeline enables accurate visual localization and motion triggering with centimeter-level precision, demonstrating the system's potential for real-world application. This foundational vision system is designed to eventually be integrated into M4's robotic arm platform, helping expand the robot's capabilities beyond locomotion and into active, autonomous manipulation. As we look ahead, this framework will be key in M4's goal of navigating and servicing hazardous electrical environments safely and intelligently.

Robotic arm for autonomous electrical panel inspection and repairs
Alexander T. Gogola
Mentors: Joel W. Burdick and Yacine Derder

Electrical panels are ubiquitous throughout any industrial plant. The goal of this project is to be able to autonomously inspect and perform repairs on electrical plants. We have designed a robotic arm to be attached to a drone. We have designed it such that it is lightweight and has a high level of precision to perform the desired tasks. Through the use of a compliant mechanism to hold the tool, we can reliably complete tasks. This research provides a general robotic arm and controller which is capable of being implemented on various mobile platforms. We show that common electrical panels tasks can be accomplished using this autonomous system.

Soft contact normal quantification using Flow Matching
Nils Jonathan Andreas Cederlund
Mentors: Joel W. Burdick and Emily A. Fourney

Robotic gasping has proven to be quite challenging, and the best approaches achieve a rate of around 90%. To improve this, we propose giving uncertainty quantification on the force normals, which in turn would allow for risk quantification of the grasp. The current state-of-the-art methods for calculating force normals provide only a single estimate without offering any risk metric. We present a Flow Matching (FM) model that samples from the distribution of normals conditioning on a point cloud. The ability to sample from the posterior distribution gives a confidence range of the force normal. The model is trained on cups, mugs, and bowls from the Shape-Net dataset. To bridge the sim-to-real gap, we use a synthetic depth image pipeline with noise models based on the ZED mini stereo camera. Our 95% confidence cone has a single angle of 18.2 deg and a mean error to the true force normal of 6.5 deg.

Numerical simulation and analysis of homogenized material properties arising from random microstructures
Rodolfo A. Ruiz
Mentors: Kaushik Bhattacharya and Harkirat Singh

Numerical simulation of material response to external stimulation (e.g., loading and heating) can be expensive because a numerical solver must resolve details on the scale of the material's microstructure when solving a problem that typically deals with much larger scales. A process called homogenization lets us estimate the solution to a differential equation of interest at a much lower computational cost by eliminating the scale of the microstructure and approximating the multi-phase system with one homogenized material property. In this project, random microstructures that consist of circles of one phase embedded in a different phase will be generated, and their homogenized properties will be solved numerically. By varying characteristics of the microstructure, such as volume fraction, number of circles, and material property contrast, we aim to identify conditions under which a limited number of statistical properties of randomly generated microstructures are sufficient to predict the statistics of the homogenized material properties accurately.

Engineering and analysis of guayule latex-algae biocomposites for sustainable material applications
Emily A. Stanton
Mentors: Chiara Daraio and Siddharth Premnath

To prevent plastic pollution and the waste generated through current biocomposite processing, we have developed a new type of biocomposite out of guayule latex and algae. We hypothesize that increasing guayule-latex volume fraction will enhance cyclic resilience while sacrificing stiffness. We compressed samples containing 0-15 wt% guayule latex with a carver press at 8 tons and tested these samples for cyclic loading and compression modulus. We also designed and manufactured a compression mold to create three-point bend test samples. We have analyzed the effect of volume fraction of latex on the cyclic loading/unloading cyclic behavior of these biocomposites and found that as latex concentrations increase by 5 wt%, the peak stresses of each sample decrease by approximately .5 MPa. The residual strain increases at a faster rate over 9 cycles as the guayule latex increases up to 10%. We will run compression tests to determine the compression modulus, the yield strength, the failure modes, and the strain at those failure modes of the biocomposites. We will also use the samples from the mold we created to perform the 3 point bend test and conduct a systematic comparison of latex-algae composites manufactured by high-pressure compression molding versus 3D printing.

Analysis of urban seismic signals from the Community Seismic Network (CSN) using K-means clustering
Pat Mutia
Mentor: Monica D. Kohler

The Community Seismic Network (CSN) consists of over 1200 stations with MEMS-based triaxial accelerometers mainly placed across the greater Los Angeles County. Due to the placement in a heavily-populated urban setting, the recorded waveform data includes significant noise from various sources such as population density, seismic site conditions (e.g. sedimentary basins with energy scattering in soft soil), anthropogenic sources (e.g. construction, traffic, and trains), and installation locations (e.g. basement versus upper building levels). To explore whether machine learning can distinguish between seismic and human-generated signals, we applied k-means clustering on pre-processed hourly noise metrics (RMS, L1 norm, and max peak ground acceleration) to assess cluster cohesion and separation across time and space. Overall, cluster separation was low but consistent patterns emerged that aligned with daily human activity cycles and school schedules. These results highlight the potential of unsupervised learning in identifying subtle recurring noise patterns in urban seismic data. Future work will focus on improved preprocessing, cyclical time encoding, and spatiotemporal visualizations to better interpret cluster structure and isolate anthropogenic influences.

Towards the design of an ultra-high vacuum compatible inductively coupled plasma (ICP) source
Annie M. Hu
Mentors: Austin J. Minnich and Mete Bayrak

Materials research is essential in enabling fabrication of cutting edge semiconductor and quantum devices, especially for nanofabrication of high-quality thin films. Plasmas, with their high density of energetic particles and reactive species, enable faster and more uniform substrate etching and/or thin film deposition. Due to the growing relevance of plasma processing for growing desirable materials in our lab, and our desire to minimize contamination and impurities in our thin films, we hope to build an ultra-high-vacuum (UHV) compatible ICP source for nitrogen plasma. In the future, we hope this technology enables growth of III-V nitrides to test in our lab. In order to create our source, we first tested using a non-UHV-compatible ICP design to see what materials we could use for the dielectric tube, what coil geometries were sufficient to strike plasma, and to observe the behavior and properties of the inductively coupled plasma under different power, pressure, and coil configurations. After understanding what is required to obtain the plasma we want and designing an appropriate grounded aperture to filter out charged species from plasma reaching the substrate, we will assemble the UHV ICP and collect measurements using a Langmuir probe to verify the functionality of the aperture.