Caltech
Student-Faculty Programs

Session I Abstracts

Toward shoreline-aware thermal water segmentation for UAVs in near-shore environments
Agnes M. Göransson
Mentors: Soon-Jo Chung, Anahita Eshghetorki, Joshua Cho, and Matthew Anderson

Autonomous systems must be able to perceive their surroundings reliably, even in visually challenging environments such as low light or fog. While RGB cameras are common, they often fail under these conditions. Thermal cameras provide a more robust signal, but labeled thermal datasets for model training are limited and costly to produce. This work explores how thermal image segmentation for water in near-shore environments can be improved using simplified class mappings and self-supervised pseudolabels, enabling the use of unlabeled data to enhance learning. Additionally, we investigate whether adding an additional label for shoreline is possible without re-annotations, while maintaining high segmentation performance. To further support perception, we explore monocular depth estimation and edge detection to provide geometric context to the segmentation model. Combined, these components create a modular and data-efficient perception pipeline that improves segmentation accuracy and robustness in challenging environments, while remaining practical for deployment in UAV-based autonomous systems.

ContractionPPO: Certified reinforcement learning via differentiable contraction layers
Narek Harutyunyan
Mentors: Soon-Jo Chung and Matthew Anderson

We present a novel framework ContractionPPO for certified robust control of legged robots by augmenting Proximal Policy Optimization (PPO) with a state-dependent contraction metric layer. This approach enables the PPO policy to not only maximize locomotion performance but also produce a contraction metric that certifies incremental exponential stability of the closed-loop system. The metric is parameterized as a Lipschitz neural network trained jointly with the policy, either in parallel or as a final head of the PPO backbone. We derive hinge-type penalty losses from a differential Lyapunov inequality, enforcing contraction constraints during training and yielding provable convergence of trajectories despite policy expressiveness and observation mapping. The overall loss combines the standard PPO objective with contraction-related penalties, positive-definiteness, conditioning, and Lipschitz regularization, leading to explicit stability margins and robustness to disturbances. We provide necessary and sufficient conditions for exponential convergence, derive upper bounds on worst-case contraction residuals, and detail the integration of metric learning into the PPO pipeline. Our hardware experiments on multiple quadruped locomotion tasks demonstrate that contraction augmented PPO enables robust, certifiably stable control, even in the presence of external disturbances.

Investigating map-free thermal navigation for aerial robots in low-light, over-water environments using imitation learning from traditional planners
Hongyi Hu
Mentors: Soon-Jo Chung, Nikhil Ranganathan, and Xingxing Zuo

This project examines whether a small robot can navigate using only passive thermal imagery when visible light is unreliable and active sensors such as lasers are undesirable. The objective is to learn a controller that converts a thermal image and a goal direction into short motion commands that keep safe clearance from obstacles. To create examples of safe behavior, we have completed an expert pipeline that builds three-dimensional maps with a laser scanner and motion sensors, computes collision-free paths with a classical planner, and time-aligns each path with the thermal video recorded in the same scene. These paired image–action examples will train a compact learning system that imitates the expert's steering and forward motion while using no active sensing at run time and avoiding reliance on visible-light cameras. Because experiments are not yet complete, the next phase expands the dataset to low-contrast scenes and water surfaces, adds brief memory over recent frames, integrates a simple safety check before each action, and evaluates performance by goal-reach rate, collision rate, and deviation from planned paths in indoor corridors and along waterfronts.

Extreme atmospheric wind sensing and quantification for adaptive control of a fixed-wing UAV
Brandon Lee
Mentors: Morteza Gharib and Xiaozhou Fan

Fixed-wing aircraft operating in turbulent environments face significant control challenges due to sudden wind gusts that exceed conventional control system capabilities. This project aims to develop and validate an experimental apparatus that evaluates an adaptive disturbance rejection control algorithm. A remote-controlled fixed-wing aircraft was modified to fit custom pitot-static probes feeding five differential pressure sensors per wing. Sensor data travels through a Teensy 4.1 microcontroller to an NVIDIA Jetson Orin Nano running Robot Operating System. A neural network model uses filtered pressure inputs to estimate instantaneous wind speed and direction. Initial verification by manual gust tests confirmed proper electronic system functionality and successful real-time data collection from pressure sensors. We will calibrate the system in a controlled wind tunnel, add servo-driven control surfaces, and perform tethered outdoor trials with fan-generated gusts.

Experimental study of new flexible solar array
Maxwell B. Braithwaite
Mentors: Sergio Pellegrino and Maria Kechri

The Caltech Space Solar Power Project aims to launch a 60m x 60m solar structure into space, harvesting solar energy, and then sending that energy back to earth via microwaves. To do this, a large deployable space structure was designed to accommodate size constraints of launch vehicles. The design of the structure includes four triangular quadrants composed of trapezoidal strips, filled with solar panels, with hinges and diagonal tapes connecting them. The diagonal tapes provide prestressing, while the hinges are there to couple the strip kinematics. Alternate hinge designs and locations were used to study the interactions between the hinges and the diagonal tapes. The information collected was used to pursue decreasing the space between the strips of solar panels. The space between the strips is unproductive, since no photo voltaic cells can be mounted there, and the amount of power lost due to the area of the gap has a significant effect. A 20% decrease in the space between strips of the solar structure has been achieved, but more is possible once a deeper understanding of the relationship between the hinges and the diagonal tapes is found that reduces contact interactions. The interactions between the strips, the new hinges, and the diagonal tapes continues to be studied to see how these factors will affect the characteristics of the deployment of the flexible structure.

High-throughput, high-strain-rate material testing and characterization via Automated Rapid Direct Impact Tester (ARDIT) for predictive modeling
Mason M. Holmes
Mentors: Guruswami Ravichandran and Aditya B. Shedge

For fiber-reinforced composites tested at strain rates above 10^2 s^{-1}, constitutive data is relatively sparse, limiting the fidelity of impact constitutive design models. This is largely due to the two conventional testing methods: Split-Hopkinson (Kolsky) bar and laser-shock having individual limiting drawbacks. Split-Hopkinson bars have lower-throughput data collection, and laser-shock methods struggle with boundary condition control. In an effort to obtain such data, and build similar constitutive design models, we have completed the construction of the Automated Rapid Direct Impact Tester (ARDIT). ARDIT is an automated gas gun capable of performing controlled planar impacts at strain rates from 10^2 -10^4 s^{-1}, while keeping track of data such as impact velocity, surface velocity, and images for Digital Image Correlation (DIC). DIC allows for a full-field measurement approach to assess deformation resulting from ARDIT impacts.

Post-completion of ARDIT, the main focus has been calibrating ARDIT's mechanical and optical systems using an isotropic, well-documented material: copper. Future phases of the project will extend testing to carbon fiber and other composites, collecting large data sets needed for machine learning. ARDIT is the fundamental experimental component for high-throughput, high-strain-rate characterization needed to feed prediction models.

Advancing ASTM E659 testing: Investigating and optimization of the current ASTM testing methods and apparatus
Jason M. Kamau
Mentors: Joseph E. Shepherd and Charline Fouchier

The autoignition temperature(AIT), defined as the minimum temperature at which a fuel ignites without an external ignition source, is a crucial parameter in assessing fire hazards across many industries. The most used methodology for determining this temperature is the ASTM-E659 standard in the United States. However, the Explosion Dynamics Laboratory, under Professor Shepherd's supervision, established that the mechanisms governing autoignition remain poorly understood, particularly in complex hydrocarbon fuels. Insights into the dispersion and ignition of such fuels under the ASTM-E659 conditions have been reached thanks to an optically accessible experimental setup. The fuel of study (hexane) was injected inside a heated test cell via a syringe pump system. Flow, concentration, and temperature data were acquired via an IR laser, photodetector, thermocouples, a high-speed camera, and LabVIEW, and analysis was done via MATLAB. The cell design was updated during the project to improve the repeatability of the tests. We evaluated injection repeatability and convection effects on concentration over a couple of tests. The calculated injected moles of fuel in the cell for all the tests were approximately 91% of the predicted moles. Stronger convection, driven by vertical temperature gradients, showed faster mixing. Finally, ignition was investigated inside the cell.