Caltech
Student-Faculty Programs

Session G Abstracts

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.

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

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.

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.

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.

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.

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.

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.

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, temperaturevariation 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.

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.

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.