Knowing When to Relax

Adaptive Control Barriers for Safe Spacecraft Proximity Operations

ABSTRACT

Autonomous spacecraft inspection and docking missions demand controllers that guarantee safety under strict thrust limits and significant uncertainty, while remaining efficient enough for real flight. Control barrier functions (CBFs) offer a principled route to safety certification by enforcing forward invariance of a prescribed safe set through a single convex quadratic program per control cycle. However, traditional CBFs are greedy, can suffer from relative degreee problems, can be invalid under uncertainities and thrust limits. Input-constrained CBFs (ICCBFs) extend traditional CBFs to respect bounded actuation and the higher-relative-degree constraints typical of proximity operations, such as approach corridors, line-of-sight cones, and keep-out zones, by introducing an inner safe set that is robust to thrust limits. In practice, however, even ICCBF formulations can be overly conservative, as they shrink the admissible set, drive up fuel consumption, and can become infeasible near constraint boundaries or under disturbances. This seminar presents two frameworks that treats the conservatism of an ICCBF as something to be learned rather than fixed by hands.

The first framework parameterizes the full hierarchy of class-$\mathcal{K}$ functions defining the ICCBF recursion and shapes it online with a meta-reinforcement-learning policy, trained across broad distributions of hidden physical parameters, state noise, and thrust uncertainty. A control margin computed efficiently via differential algebra allows the continuous-time barrier to be enforced safely on time-sampled, zero-order-hold systems without resorting to intensive grid-based estimation. The result is a safety filter that learns to relax toward fuel efficiency where it can, and to tighten where it must prioritize safety. A systematic study was also conducted to determine which network architectures and learning algorithms can best learn the class-$\mathcal{K}$ hierarchy, comparing recurrent (Long-Short Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures) and selective state-space (Mamba) backbones under both on-policy (Proximal Policy Optimization (PPO)) and off-policy (Soft-Actor-Critic (SAC)) training.

The second framework develops a unified two-stage CBF design where RL is used to train two separate networks. The first network acts similarly to that of the earlier framework, exploiting class-$\mathcal{K}$ functions to develop non-greedy CBFs. The second network's purpose is to expand the inner safe set of ICCBFs, as they conservatively abandon states that are in the original safe set. The second network trains a neural function that returns these states to the inner safe set, effectively expanding the viability kernel. At runtime the framework simply selects the appropriate stage based on the state and solves a single lightweight QP under zero-order hold, such that the analytical guarantees and real-time efficiency of the CBF–QP structure are retained, while the effective operating region is enlarged toward the true viability kernel.

For both frameworks, the results show that the learned filters maintain safety while substantially reducing fuel use relative to fixed-parameter ICCBFs, with recurrent policies in Framework 1 proving especially effective when more hidden parameters are at play. The systematic study extends the test cases to adversarial settings, in which a non-cooperative target actively evades the chaser by altering its rotation rate, or denies sensor coverage by translating away — forcing the policy to infer latent adversarial intent from trajectory history and modulate its barrier shaping accordingly. This study shows that state space architectures outperform traditional LSTM networks when learning CBFs. Lastly, the two-stage formulation in Framework 2 reduces the median fuel consumption relative to the ICCBF baseline while recovering enough abandoned states to raise the fraction of trajectories that remain safe.

SPEAKER BIO

Minduli Wijayatunga

Incoming Assistant Professor – University of Illinois, Urbana-Champaign

Minduli Wijayatunga is an adjunct and incoming Assistant Professor at the University of Illinois Urbana–Champaign, and is currently also a Visiting Scientist at MIT. Her research focuses on guidance, navigation, and control for spacecraft proximity operations, with particular emphasis on safety and machine learning, and spans control barrier functions, reinforcement learning, indirect methods, and convex-optimization-based trajectory design. She is an Amelia Earhart Fellow and a Forbes 30 Under 30 honoree. She received her Ph.D. in Astrodynamics from the University of Auckland, New Zealand, and her B.Sc. and B.Eng. in Aerospace Engineering and Science from the University of Sydney, Australia.