Paper ID: 75966
student

IAF ASTRODYNAMICS SYMPOSIUM (C1)
Guidance, Navigation and Control (1) (3)

Author: Mr. Omkar Mulekar
University of Florida, United States, omkar.mulekar@ufl.edu

Dr. Riccardo Bevilacqua
Embry-Riddle Aeronautical University, United States, bevilacr@erau.edu
Dr. Hancheol Cho
Embry-Riddle Astronautical University, United States, CHOH15@erau.edu

STABILITY OF DEEP NEURAL NETWORKS FOR FEEDBACK-OPTIMAL PINPOINT LANDINGS

Abstract

Modern computing has motivated optimization of highly parametrized regression techniques like neural networks, allowing them to automate tasks like handwriting recognition, speech recognition, vehicle navigation, and autonomous control. In the realm of spacecraft guidance, navigation, and control (GNC), neural networks have been trained via reinforcement learning (RL) and imitation learning (IL) techniques to perform autonomous feedback control for a variety of tasks like hovering, orbit transfers, and pinpoint landings. RL and IL have shown some success at learning optimal feedback control for aerospace vehicles, though the trained policy does not undergo analytical stability assessments. While RL allows policies to learn in the environment, reward functions are not straight-forward to design for optimal control problems since a maximized reward does not translate directly to a minimized cost in the optimal control sense. IL seeks to directly leverage "expert" demonstrations (e.g. numerically generated open-loop optimal trajectories) to train a feedback-optimal policy. IL has shown success in low and high-fidelity pinpoint landing formulations, though stability was shown via demonstration which is not as strong as an analytical assessment. Some studies have investigated stability guarantees to the linear-quadratic-regulator (LQR) problem extended to nonlinear dynamics by balancing trained policies in states far from equilibria with the LQR in approximately-correct linearized regions of the state-space about the equilibria. Despite the recent progress autonomous control, neural networks still lack analytical stability guarantees and assessments for nonlinear optimal control problems. This study addresses the stability issue by training a policy to perform pinpoint landings and formulating the IL problem as an inequality-constrained optimization problem. The network is to be trained on optimal example trajectories, where the network loss function is constrained to a Lyapunov inequality. Then, assessment on this stability is shown by finding critical points of the Lyapunov derivative, evaluating the Hessian matrix and showing that maximums of the Lyapunov derivative in a subspace of the state-space are negative. The developed method will be investigated for a variety of lander formulations from low-fidelity 3-degree-of-freedom landers to high-fidelity 6-degree-of-freedom landers. The developed method will show that controller stability in the Lyapunov sense can be assessed across the distribution of states in the training data, and that the IL problem can be posed to enforce Lyapunov stability during training. This technique allows for feedback-optimal control of pinpoint landings with stability assessment capabilities.