

IAF ASTRODYNAMICS SYMPOSIUM (C1)
Interactive Presentations - IAF ASTRODYNAMICS SYMPOSIUM (IP)

Author: Mr. FNU Vedant
University of Illinois, United States, vedant2@illinois.edu

Prof. James Allison
University of Illinois, United States, jtalliso@illinois.edu

Prof. Matthew West
University of Illinois, United States, mwest@illinois.edu

Dr. Alexander Ghosh
University of Illinois, United States, aghosh2@illinois.edu

REINFORCEMENT LEARNING FOR SPACECRAFT ATTITUDE CONTROL

Abstract

Reinforcement learning (RL) has recently shown promise in solving difficult numerical problems and has also discovered non-intuitive solutions to existing problems. We investigate the ability of a general reinforcement learning agent to find an optimal control strategy for the spacecraft attitude control problem.

We consider the general ADCS (Attitude Determination and Control System) problem with full actuation, but with saturation constraints on the applied torques due to limits from attitude thrusters. The reward function for the controller includes terms for the attitude error and applied torque, and we demonstrate the inclusion of further terms including nonlinear functions of the system state. Once the general ADCS problem is solved, a candidate problem with higher fidelity torque actuator model is solved. Specifically, models for torques produced are based on reaction wheels/magnetorquers; these provide more complete constraints on control authority.

The agent is trained using the Proximal Policy Optimization (PPO) reinforcement learning method to obtain an attitude control policy. To ensure robustness, the inertia of the satellite is unknown to the control agent and is randomized for each simulation rollout. To achieve efficient learning, the agent is trained using curriculum learning and Hindsight Experience Replay (HER).

We compare the reinforcement-learning controller to a QRF (quaternion rate feedback) attitude controller (a well-established state feedback control strategy), and we investigate the nominal performance and robustness with respect to uncertainty in system dynamics. Our results suggest that reinforcement learning can deliver superior attitude control schemes in some regimes, and we discuss the tradeoffs for ADCS design.