Aerospace Europe Conference 2023
Joint 10th EUCASS – 9th CEAS Conference

Abstract #XXX *(to be filled by the organizers)*
Preferred Topics: FDGNCAV
Corresponding author: Jan Luca Loettgen
e-mail of corresponding author: j.loettgen.1@research.gla.ac.uk
Type: Oral
Status of corresponding author: Student
For student corresponding author: student member of one of the following: **RAeS**

Title

Deep reinforcement learning for spacecraft attitude tracking manoeuvres

Authors

Jan Loettgen [1], Matteo Ceriotti [2], Gerardo Aragon-Camarasa[3], Kevin Worrall[4]*

** Corresponding author*
[1] James Watt School of Engineering, University of Glasgow, G12 8QQ, Glasgow, United Kingdom,
j.loettgen.1@reserach.gla.ac.uk
[2] James Watt School of Engineering, University of Glasgow, G12 8QQ, Glasgow, United Kingdom,
matteo.ceriotti@glasgow.ac.uk
[3] School of Computing Science, University of Glasgow, G12 8RZ, Glasgow, United Kingdom,
gerardo.aragoncamarasa@glasgow.ac.uk
[4] James Watt School of Engineering, University of Glasgow, G12 8QQ, Glasgow, United Kingdom,
kevin.worrall@glasgow.ac.uk

Abstract

Deep reinforcement learning has shown promise in providing an end-to-end attitude control strategy that can be utilised for different spacecraft without control parameter tuning. Furthermore, by modifying the reward function used to train the reinforcement learning agent, the same controller can be trained to complete either minimum time or minimum control effort manoeuvres. Existing research into the applications of deep reinforcement learning to satellite attitude control has almost exclusively focused on attitude regulation manoeuvres, where the target orientation is stationary in the inertial reference frame [1-3]. When the target orientation is varying with time, then the state vector loses the Markov Property with the Markov Decision Process that the agent is learning to solve becoming only partially observable. A common method for addressing partial observability is frame stacking, where the agent is provided with a specified number of past observation vectors, effectively providing it with perfect short-term memory. A second method uses Long Short-Term Memory (LSTM) layers in both the actor and critic networks which provide the agent with imperfect short- and long-term memory [4]. In this work, a deep reinforcement learning agent is trained to control the orientation of a 6U-CubeSat, tracking a time series along Bezier curves. An actor-critic Proximal Policy Optimisation deep reinforcement learning algorithm is used to train a continuos action space controller in the Basilisk astrodynamics simulator. Two further agents are trained that use frame-stacking and LSTM layers to account for partial observability. The performance of all three controllers is compared against a standard tuned PD controller in terms of the mean pointing error.

References

[1] Vedant JT. Reinforcement Learning for Spacecraft Attitude Control. In70th International Astronautical Congress 2019 Oct.

[2] Elkins J, Sood R, Rumpf C. Adaptive continuous control of spacecraft attitude using deep reinforcement learning. In 2020 AAS/AIAA Astrodynamics Specialist Conference 2020 (pp. 420-475). Reston, VA: AIAA.
[3] Elkins J, Sood R, Rumpf C. Autonomous spacecraft attitude control using deep reinforcement learning. In 71st International Astronautical Congress (IAC) 2020 Oct (Vol. 2020).
[4] Meng L, Gorbet R, Kulić D. Memory-based deep reinforcement learning for pomdps. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2021 Feb 25 (pp. 5619-5626). IEEE.