Proceedings of the 1st international Symposium on AI and Fluid Mechanics
Abstract No S15 A6

A hybrid model-based/model-free approach for flight control of flapping wing drones using reinforcement twinning

Romain Poletti[1, 2], Lorenzo Schena[1, 3], Lilla Koloszar[1], Joris Degroote[2]
Miguel Alfonso Mendez[1]

[1] *the Von Karman Institute (VKI) for Fluid Dynamics*

[2] *Ghent Universiteit*

[3] *Vrije Universiteit Brussel (VUB)*

Abstract

Controlling flapping-wing drones requires versatile flight controllers that handle their time-varying, non-linear, and underactuated dynamics from incomplete and noisy sensor data. Model-based methods struggle with accurate modeling, while model-free approaches falter in efficiently navigating very high-dimensional and nonlinear control objective landscapes. This article presents a novel hybrid model-free/model-based approach to flight control based on the recently proposed reinforcement twinning algorithm [1]. The model-based (MB) approach relies on an adjoint formulation using a time-invariant model, continuously identified from live trajectories, while the model-free (MF) approach relies on reinforcement learning. The two control loops interact and exchange learning feedback by interacting with (1) the real environment, (2) an adaptive model of the environment acting as a digital twin and (3) a referee. While both MB and MF agents are continuously trained with the digital twin, the referee selects only one of these to interact with the real environment. This enables collaboration between agents via transfer learning, imitation learning and experience sharing. The referee is driven by performances on the digital twin and a model-to-environment trust ratio that weights control performances in the virtual environment with model accuracy. The algorithm is evaluated for controlling the longitudinal dynamics of a flapping-wing drone. The real environment is simulated using a non-linear, time-varying dynamical system. The hybrid control learning approach is tested with three types of initialization of the adaptive model: (1) offline identification using previously available data, (2) random initialization with full online identification, and (3) offline pre-training with an estimation bias, followed by online adaptation. In all three scenarios, the proposed hybrid learning approach demonstrates superior performance compared to purely model-free and model-based methods.

Keywords:

Model-based Reinforcement Learning, Adjoint-based Assimilation, Transfer Learning, Imitation Learning, Experience Sharing

References

[1] Schena, L., Marques, P. A., Poletti, R., Ahizi, S., Van den Berghe, J., & Mendez, M. A. (2024). Reinforcement twinning: From digital twins to model-based reinforcement learning. Journal of Computational Science, 82, 102421.