Proceedings of the 1st international Symposium on AI and Fluid Mechanics
Abstract No S15 A2

Reinforcement Twinning: from digital twins to model-based reinforcement learning

Lorenzo Schena[1,2], **Pedro Marques**[1,3], **Roman Poletti**[1,4], **Samuel Ahizi**[1,5], **Jan van den Berghe**[1,6], **Miguel A. Mendez**[1]

[1] *the Von Karman Institute (VKI) for Fluid Dynamics,* [2] *Vrije Universiteit Brussel (VUB),* [3] *Université Libre de Bruxelles (ULB),* [4] *University of Ghent,* [5] *Universidad Carlos III de Madrid,* [6] *Université Catholique de Louvain*

Corresponding author: Miguel Alfonso Mendez (mendez@vki.ac.be)

As machine learning and data-driven approaches are increasingly used to tackle control problems, their direct deployment onto real-world systems is still hindered by their lack of guarantees and high sample complexity, which often clash against experimental physical constraints. This leads to a dichotomy between adaptability and robustness, in which the model-free capabilities of learning control laws directly from measurement must be traded off against the stability and transparency offered by physically rooted model-based methods. This work seeks to combine the best of these two paradigms, proposing a framework that we refer to as *Reinforcement Twinning* (RT, [1]).

The RT architecture is composed of two main blocks, shown in Fig. 1. Namely, (1) an adaptive digital twin of the plant, and (2) a controller that intertwines model-based and model-free methods. The digital twin relies on a physics-based model and a parametric closure law, that encloses unknown or uncertain properties of the system, that we tune via an adjoint-based data assimilation strategy. The control-design loop evolves simultaneously a model-based 'optimal' controller and a model-free reinforcement learning component, selecting the best one to drive the next interaction with the real system. This triggers a virtuous loop, in which the deployed controller effectively acts as a 'master' for the other one, which can now learn from expert demonstrations.

We test the proposed RT framework on three different control tasks: (1) the control of a wind turbine subject to turbulent winds, (2) the trajectory control of flapping-wing micro air vehicles (FWMAVs) exposed to wind gusts, and (3) the mitigation of thermal loads in the management of cryogenic storage tanks. The adjoint-based twinning is remarkably efficient, converging in a few iterations for all analysed problems. The training of the control agents proves the benefit of the proposed interaction, which leads to a reciprocal faster convergence. These encouraging results open the path towards deploying the RT framework on real systems.

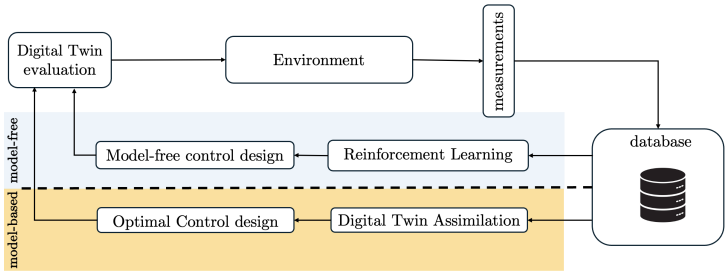

Figure 1: Architecture of the Reinforcement Twinning (RT) framework.

References

[1] Lorenzo Schena, Pedro A Marques, Romain Poletti, Samuel Ahizi, Jan Van den Berghe, and Miguel A Mendez. Reinforcement twinning: From digital twins to model-based reinforcement learning. *Journal of Computational Science*, 82:102421, 2024.