International Forum on Aeroelasticity and Structural Dynamics
IFASD 2026
17-21 June 2026, Goettingen, Germany

CO-DESIGN OF AEROELASTIC SYSTEMS WITH DEEP REINFORCEMENT LEARNING

Yao Cheng Li[1] and Urban Fasel

[1]Imperial College London
Exhibition Road, SW72AZ London, United Kingdom
paul.li20@imperial.ac.uk

ABSTRACT

1 INTRODUCTION

Co-design is a nested or simultaneous optimisation process to optimise both the system and controller. By treating these often highly coupled elements as one, it can reach an optimal solution that a sequential optimisation approach might not achieve [1–3]. One such example is an aeroelastic system navigating through a thermic soaring environment that exploits atmospheric phenomena to extract energy [4–6]. Certain phases of flight require the execution of agile manoeuvres to exploit highly localised thermals, which contradicts directly with an optimal high aspect ratio wing design with high roll inertia and damping. Furthermore, thermal soaring presents an additional challenge of locating thermals and path planning. Deep reinforcement learning has demonstrated remarkable performance in navigating partially observed environments [7–12], and has recently been applied to co-design [13–17]. This motivates the development of a deep reinforcement learning co-design framework for a fixed-wing glider in a thermal soaring environment, shown in Figure 1. In this work, we present 1) a low-fidelity flight simulation with a fully parameterisable elastic high aspect ratio fixed-wing glider model, 2) develop control strategies using deep reinforcement learning for highly flexible aircraft, and 3) integrate data-driven co-design methods in optimising a glider and its control in a unifying model-free framework.

2 PROBLEM STATEMENT

A key component of the co-design training process is formulating a well-defined dynamical environment that is also computationally efficient to facilitate reinforcement learning. Here, we developed a flexible aircraft simulation environment in Mujoco [18] by adapting a rigid multi-body implementation of a flexible wing [19]. Mujoco, which stands for multi-joint dynamics with contact, is a physics engine designed to facilitate model and reinforcement learning-based control developments, widely used in the reinforcement learning research community [20]. Many physics extensions have already been developed in Mujoco, such as flapping flight for the simulation of a fruit fly [21]. While many aeroelastic and flight dynamics solvers are available, our goal is to facilitate the use of reinforcement learning for the design and control of flexible aircraft that exhibit aeroelastic phenomena. Therefore, we extend Mujoco and exploit its existing functionalities.

In our proposed rigid panel wing framework, shown in Figures 2, 3 and 4, each wing section is represented by a body with prescribed inertia properties and structural elasticity. The elasticity

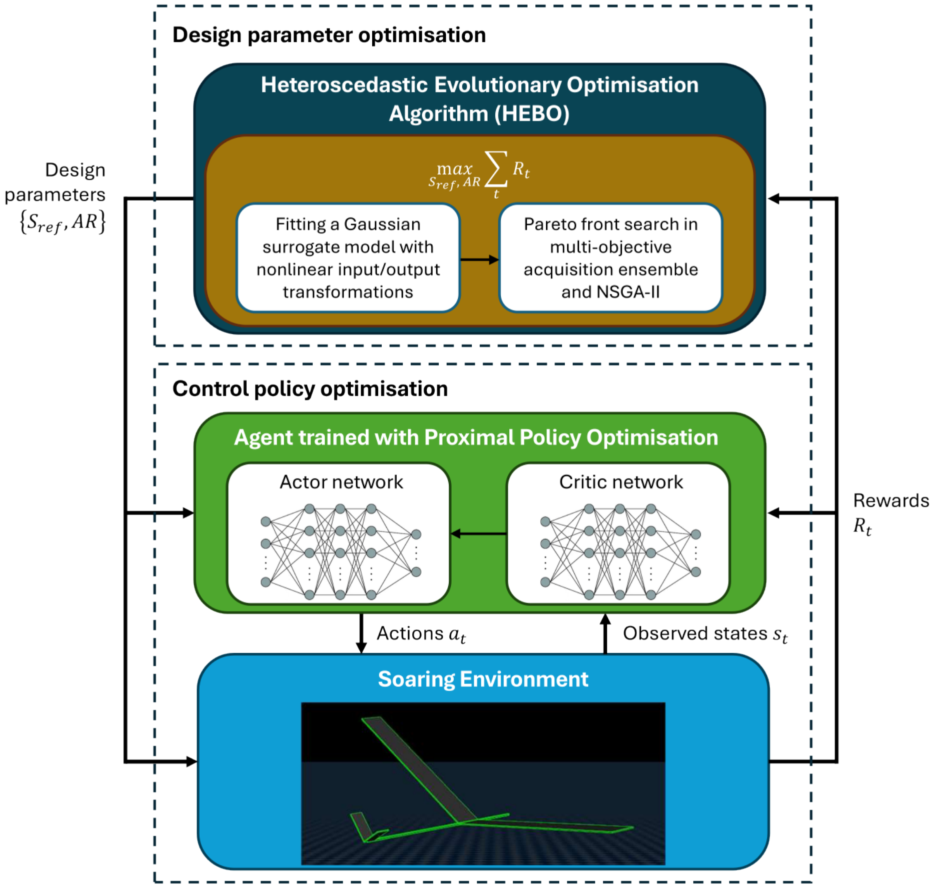

Figure 1: Overview of the co-design process.

is modelled using two torsional springs with linear stiffnesses to represent torsional and out-of-plane bending. A preview of the implementation is shown in Figure 2, where a straight, non-tapered wing is discretised into discrete strip sections. Viterna's theory was implemented to model the nonlinear lift characteristic [22].

The soaring environment consists of multiple thermal and sink columns, modelled as radial basis functions, as shown in Figure 6. Our nested co-design formulation, depicted in Figure 1, jointly optimises the wing area S_{ref}, aspect ratio AR, and its control policy to both control the aircraft and navigate the environment. The objective is to maximise the total reward gained by the agent by maximising the number of mission waypoints reached whilst avoiding ground collision. The agent observes its own state variables, local environment variables, and mission-related parameters, as well as the design sampled from the design space from the outer loop. It

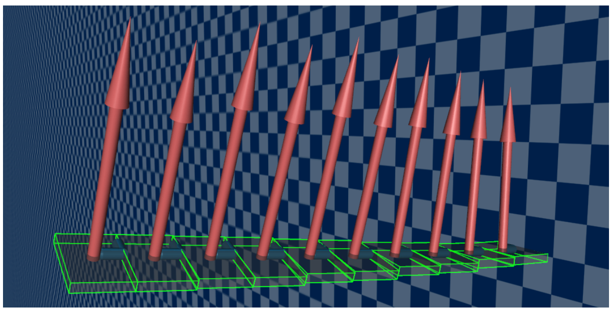

Figure 2: Example discretisation of a clamped port-side wing in Mujoco.

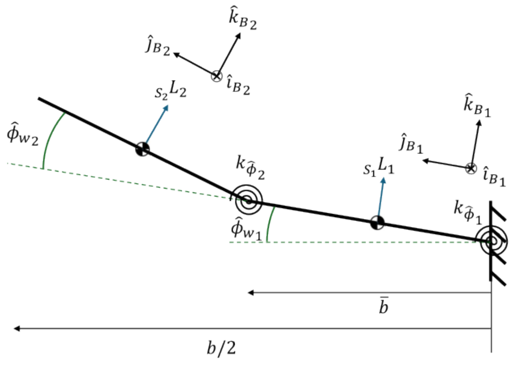

Figure 3: Rear view of a 2 section discretisation of the port side wing.

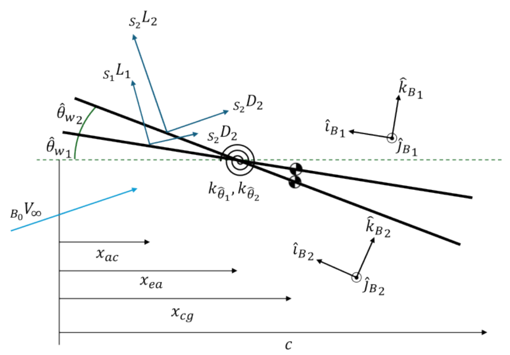

Figure 4: Port side view of a 2 section discretisation.

then interacts with the environment through direct control of the trailing edge flap deflections of its wing sections, as prescribed by its control policy. It receives a tuned reward for its actions, such as reaching a waypoint, which is then recorded by both inner and outer optimisation loops.

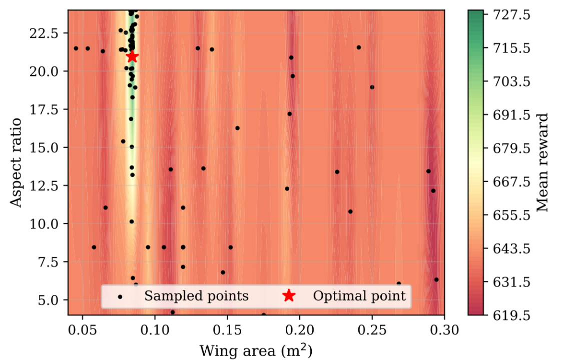

Figure 5: Co-design reward landscape.

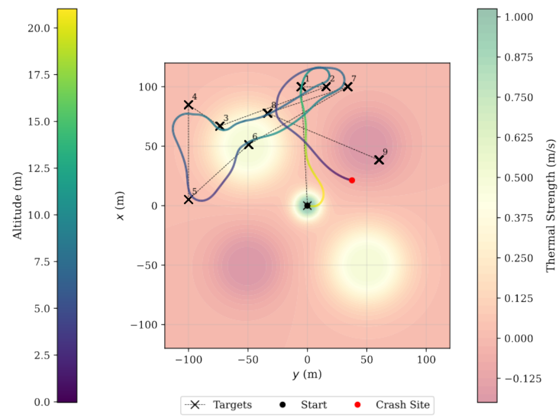

Figure 6: Flight path trajectory undertaken by the advanced agent during evaluation.

The candidate optimal design produced by the proposed formulation is benchmarked against a baseline aircraft optimised solely for its glide ratio. Preliminary studies show that the co-designed aircraft with a 65% smaller wing area and 22% larger aspect ratio has achieved 70% higher mean reward compared to the baseline aircraft. Its reward landscape is shown in Figure 5, and an example trajectory is shown in Figure 6.

3 FUTURE WORK

For the full paper submission, we intend to

- Verify the low-fidelity aeroelastic framework and quantify its errors.
- Robustify the optimisation process through design of the agent's observations, actions and its reward structure.
- Explore alternate aeroelastic formulation such as unsteady vortex lattice methods.

4 REFERENCES

[1] Garcia-Sanz, M. (2019). Control co-design: an engineering game changer. *Advanced Control for Applications: Engineering and Industrial Systems*, 1(1), e18. ISSN 2578-0727.

[2] Fathy, H. K., Reyer, J. A., Papalambros, P. Y., et al. On the coupling between the plant and controller optimization problems. In *Proceedings of the 2001 American Control Conference.(Cat. No. 01CH37148)*, vol. 3. IEEE. ISBN 0780364953, pp. 1864–1869.

[3] Herber, D. R. and Allison, J. T. (2019). Nested and simultaneous solution strategies for general combined plant and control design problems. *Journal of Mechanical Design*, 141(1), 011402. ISSN 1050-0472.

[4] Zhao, Y. J. (2004). Optimal patterns of glider dynamic soaring. *Optimal control applications and methods*, 25(2), 67–89. ISSN 0143-2087.

[5] Richardson, P. L. (2011). How do albatrosses fly around the world without flapping their wings? *Progress in Oceanography*, 88(1-4), 46–58. ISSN 0079-6611.

[6] Pennycuick, C. J. (1998). Field observations of thermals and thermal streets, and the theory of cross-country soaring flight. *Journal of Avian Biology*, 33–43. ISSN 0908-8857.

[7] Silver, D., Huang, A., Maddison, C. J., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587), 484–489. ISSN 1476-4687.

[8] Vinyals, O., Babuschkin, I., Czarnecki, W. M., et al. (2019). Grandmaster level in starcraft ii using multi-agent reinforcement learning. *nature*, 575(7782), 350–354. ISSN 1476-4687.

[9] Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

[10] Jumper, J., Evans, R., Pritzel, A., et al. (2021). Highly accurate protein structure prediction with alphafold. *nature*, 596(7873), 583–589. ISSN 1476-4687.

[11] Abramson, J., Adler, J., Dunger, J., et al. (2024). Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016), 493–500. ISSN 0028-0836.

[12] Muir, M. J., Rolston, S., Shehbaz, Z., et al. Multi-objective, multi-agent based approaches for predication of sustainable aviation rollout. In *AIAA SCITECH 2024 Forum*. p. 2015.

[13] Bergonti, F., Nava, G., Wüest, V., et al. Co-design optimisation of morphing topology and control of winged drones. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. ISBN 9798350384574, pp. 8679–8685.

[14] Buda, A. A., Grover, D., Armanini, S., et al. (2024). Co-design optimization of adaptive structures using deep reinforcement learning. URL: `https://eccomas2024.org/event/contribution/1bcf6c84-9683-11ee-8a2d-000c29ddfc0c`.

[15] Sadat, E., Lotfalian Saremi, M., and Bayrak, A. E. A two-timescale reinforcement learning approach for control co-design problems. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 87301. American Society of Mechanical Engineers. ISBN 0791887308, p. V03AT03A003.

[16] Schaff, C. and Walter, M. R. (2022). N-limb: Neural limb optimization for efficient morphological design. *arXiv preprint arXiv:2207.11773*.

[17] Schaff, C., Yunis, D., Chakrabarti, A., et al. Jointly learning to construct and control agents using deep reinforcement learning. In *2019 international conference on robotics and automation (ICRA)*. IEEE. ISBN 153866027X, pp. 9798–9805.

[18] Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE. ISBN 1467317365, pp. 5026–5033.

[19] Havaza, V. S. and Nikitin, R. (2024). Mathematical method for prediction aeroelastic phenomena and multidisciplinary optimization lifting surfaces of flight vehicle at preliminary stage design.

[20] Kaup, M., Wolff, C., Hwang, H., et al. (2024). A review of nine physics engines for reinforcement learning research. *arXiv preprint arXiv:2407.08590*.

[21] Vaxenburg, R., Siwanowicz, I., Merel, J., et al. (2024). Whole-body simulation of realistic fruit fly locomotion with deep reinforcement learning. biorxiv.

[22] Tangler, J. and Kocurek, D. Wind turbine post-stall airfoil performance characteristics guidelines for blade-element momentum methods. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*. p. 591.