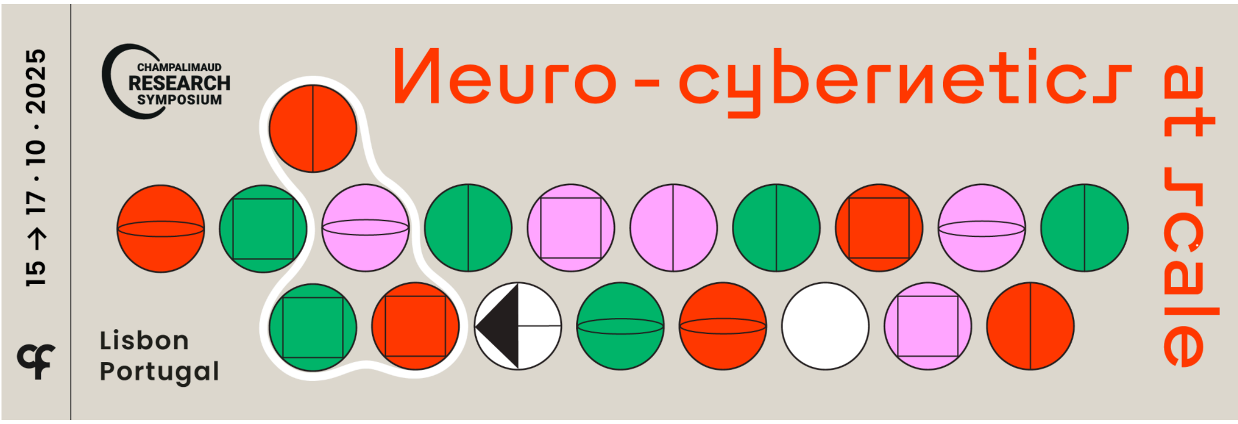

Deep reinforcement learning to mimic neuromechanical control: realistic locomotion in an embodied fly

Elliott T. T. Abe 1, Charles Zhang 2, Raveena Chhibber 1, Grant M. Chou 1, Dang P. Q. Truong 1, Bence P. Ölveczky 2, Nathan Sniadecki 1, John C. Tuthill 1, Talmo D. Pereira 3, Bing W. Brunton 1

1. University of Washington, Seattle, USA
2. Harvard University, Boston, USA
3 Salk Institute for Biological Studies, San Diego, USA

Most locomoting animals can maintain motor control after limb injury or amputation, where coordinating movements of a drastically altered body requires dynamic interactions between the nervous system, the biomechanics of the body, and the physical environment. Until recently, it has been challenging to link neural and biomechanical models to investigate adaptive motor recovery because it requires coupling neural control in closed loop with the environment. Further, while previous approaches have been able to simulate walking in biomechanically realistic bodies, the joint kinematics and ground reaction forces have been unrealistic or unvalidated. In this project, we developed and trained an agent with deep reinforcement learning (DRL) to imitate real Drosophila walking and flying using a biomechanically realistic fly body model in the physics simulator MuJoCo [1]. We implemented a variational encoder-decoder architecture [2], to build an interpretable latent space for motor control. For training data, we used high-fidelity 3D tracking data to imitate behaviors. We show that our model closely resembles real fly locomotion while reproducing accurate movement dynamics (i.e. ground reaction forces in walking). Historically, measuring forces produced by such small animals has been impossible. We validate our walking MuJoCo model with the first-of-their-kind measurements of ground reaction forces in freely walking fruit flies, demonstrating that simulated ground contact forces during walking closely match experimental measurements. Additionally, we implemented a biologically realistic aerodynamics model [3] to simulate the forces experienced by flying insects. To train our agent on higher-level tasks while maintaining realistic movements, we built a high-level controller by freezing the decoder, which translates the latent space into realistic movement, and trained a new encoder with a task driven computational goal. Using this pipeline, we can simulate in real time the full neuromechanical control in closed loop with the environment during locomotion to test and generate hypotheses.

More broadly, this work is a key step in using embodied agents to understand the neural mechanisms controlling robust movement with a dynamically changing body and environment.

embodied agents, neuromechanical control, reinforcement learning,