

## Learning for Control



Figure 1.6: Muscle-based robotic arm serving as a testbed for Learning for Control. While it offers unique possibilities in terms of high accelerations, extreme speeds, and variable stiffness actuation, classical control methods are unable to unlock these abilities.

Control of complex plants or systems, especially robots actuated by pneumatic artificial muscles, is a challenging task due to nonlinearities, hysteresis effects, massive actuator delay and unobservable dependencies such as temperature. Such plants and robots require much more from the control than classical methods can deliver. Therefore, we aim to develop novel methods for learning control that can deal with high-speed dynamics and muscular actuation.

Highly dynamic tasks that require large accelerations and precise tracking usually rely on accurate models and/or high gain feedback. While kinematic optimization allows for efficient representation and online generation of hitting trajectories, learning to track such dynamic movements with inaccurate models remains an open problem. To achieve accurate tracking for such tasks in a stable and efficient way, we have proposed a series of novel adaptive Iterative Learning Control (ILC) algorithms that are implemented efficiently and enable caution during learning [6].

Muscular systems offer many beneficial prop-

erties to achieve human-comparable performance in uncertain and fast-changing tasks [244]. For example, muscles are backdrivable and provide variable stiffness while offering high forces to reach high accelerations. Nevertheless, these advantages come at a high price as such robots defy classical approaches for control. We have built a muscular robot system to study how to accurately control musculoskeletal robots by learning control. We have shown how probabilistic forward dynamics models can be employed to control complex musculoskeletal robots on an antagonistic pair of pneumatic artificial muscles using only one-step-ahead predictions of the forward model and incorporating model uncertainty.

In addition, we have continued to work on reinforcement learning problems, at the intersection of control and machine learning. We have extended several approaches in reinforcement learning for continuous control (NAF, Q-Prop, IPG, TDM) to handle function approximations with significantly improved sample efficiency [131, 148, 174, 179, 219]. In [177], we have shown that our approach scales to learning a door opening task. Aside from fundamental algorithmic problems such as sample efficiency and stability, we also proposed algorithms that enable learning on real-world robots with less human interventions during learning. In [147], we propose the Leave No Trace (LNT) algorithm that significantly reduced the number of hard resets required during learning, paving a path toward autonomous, reset-free learning in real environments. Lastly, we made a contribution to the field of hierarchical reinforcement learning with the HIRO algorithm [106], a scalable off-policy HRL algorithm with substantially improved sample efficiency on difficult continuous control benchmarks over previous methods.

More information: <https://ei.is.mpg.de/project/learning-4-control>