Accelerating physical systems with Imagination Models

Arindam Saha[1], Aaron D. Tranter[1], Ben C. Buchler[1] and Ping K. Lam [1,2]
[1] Center for Quantum Computation and Communication Technology, Department of Quantum Science and Technology, Australian National University, Canberra, ACT 2601
[2] A*STAR Quantum Innovation Centre (Q.InC), Institute of Materials Research and Engineering (IMRE), Agency for Science, Technology and Research (A*STAR)

Machine Learning (ML) has found its way into almost every sphere of our lives, automating and optimising tasks that once took hours of manual labour. Deep Learning has been shown to outperform solutions developed by humans in several physical experiments of complex dynamics. Emerging Deep Reinforcement Learning (RL) techniques have shown massive success in control, decision-making and planning tasks, and learning optimal policies to achieve specific goals, without prior knowledge of the system [1]. However, given the long time to train, most RL algorithms are notoriously hard to train optimally. This usually isn't a big challenge when dealing with tasks virtually, where parallelisation can be applied. However, for applying to real-world systems, the RL agent must be as sample-efficient as possible. Model-based RL has shown great potential in sample efficiency and asymptotic performance [2].

In this work, we implement such an algorithm in real-world experiments. We take a black-box approach to the problem being as agnostic as possible about the underlying details of the experiment, making it a more universal approach. Having access to only the observations from the experiment, and the actions taken, we model the hidden dynamics by training neural networks, that evolve with time, thus making them highly scalable. Equipped with such models, the learning agent, can predict future transitions and develop optimal policies through imagination, making the whole process extremely efficient [3]. We demonstrate the performance of our algorithm on the task of automating the alignment of an optical cavity and tuning of a qubit system based on silicon quantum dots [4] and show its capability in running the systems optimally and sustaining their coherence over time. Our method outperforms standard optimisation and control techniques including RL ones, and establishes its efficacy at generalising across multiple platforms.

References

[1] R. S. Sutton and A.G. Barto., Reinforcement Learning: An Introduction, Vol. 9., (1998)

[2] Theophane Weber et.al., Imagination-augmented agents for deep reinforcement learning, (2017)

[3] D. Hafner, T. Lillicrap, J. Ba, M. Norouzi, Dream to Control: Learning behaviours by latent imagination, ICLR (2020)

[4] Jonathan Y. Huang et. al., High-fidelity spin qubit operation and algorithmic initialization above 1 K, Nature, 627(8005):772–777 (2024)