FM 23: Quantum & Information Science: Neural Networks, Machine Learning, and Artificial Intelligence I

Time: Monday 16:30–18:30 Location: 3043

Invited Talk FM 23.1 Mon 16:30 3043

Learning to violate Bell inequality with reinforcement learning — •Alexey Melnikov, Pavel Sekatski, and Nicolas Sangouard — Department of Physics, University of Basel

Quantum experiments push the envelope of our understanding of fundamental concepts in quantum physics. The designing of modern quantum experiments is difficult and often clashes with human intuition. In my talk, I will address the question of whether a reinforcement learning agent can propose novel quantum experiments. In our works, we answer this question in the affirmative in the context of quantum optics experiments, although our techniques are more generally applicable. I will talk about reinforcement learning and demonstrate how the projective simulation model can be used to design quantum experiments and discover experimental techniques by considering two examples. In the first example, a reinforcement learning agent learns to create high-dimensional entangled multiphoton states [1]. In the second example, our reinforcement learning agent learns to design quantum experiments in which photon pairs violate a Bell inequality. As a result of this learning process, the agent finds several optical setups with high CHSH values for various detection efficiencies, which is an important step towards realistic device-independent quantum cryptography. Our findings highlight the possibility that machine learning could have a significantly more creative role in future quantum experiments.

[1] A.A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H.J. Briegel. Proc. Natl. Acad. Sci. U.S.A., 115(6):1221, 2018

Invited Talk FM 23.2 Mon 17:00 3043

Quantum policy gradient methods for reinforcement learning — •Sofiene Jerbi[1], Hans Briegel[1,2], and Vedran Dunjko[3] — [1]University of Innsbruck, Innsbruck, Austria — [2]University of Konstanz, Konstanz, Germany — [3]University of Leiden, Leiden, Netherlands

Recent advances in quantum reinforcement learning have been considering a fully quantum learning scenario, where agent, environment and their interaction are granted quantum mechanical abilities. Such scenario is of particular interest for simulatable environments (e.g., fully specified Markov Decision Processes (MDPs)) with large state and action spaces. Indeed, in this case, finding the agent's optimal policy by solving the MDP becomes computationally intractable, which constraints the agent to learn the optimal policy through (quantum) interaction with the environment. The questions that naturally arise are then: what advantage in learning efficiency can we get from a quantum interaction? How can we exploit a quantum interaction to extract useful information for a classical agent? In this work, we tackle these questions by defining a quantization of so-called policy gradient methods. These latter turn the reinforcement learning task into a direct optimization problem on a set of parameters characterizing the agent's policy (e.g., the weights of a neural network) "solvable" with gradient ascent on these parameters. In our quantization, we explore how a quantum interaction can speed-up the computation of the gradients required for these methods.

FM 23.3 Mon 17:30 3043

Machine learning for quantum chemistry with quantum computers — •Tomislav Piskor — HQS Quantum Simulations, Karlsruhe, Germany

Simulating chemical systems is a major field of interest not only for the pharma and chemistry, but also for the automotive industry. One such example is the simulation of functional groups of a large molecule or proteine, which can be useful for the development of new medicine. In order to get the exact ground state, we might use quantum computers in the future. However, every call to a quantum computer will be relatively expensive, making high-throughput simulations with quantum computers unfeasible.

To bypass this, a few single point calculations are determined with an expensive method and then extended to more conformations with, e.g., machine learning methods. The less time-consuming method of choice is density functional theory (DFT). Our approach is to take a hybrid functional, such as B3LYP, which consists of three exchange and two correlation functionals. Each of these functionals has a certain weight which can be modified. Using, for example, a root finding optimizer the functional parameters are optimized in such a way that the energy and the corresponding nuclear gradients of the time-consuming method match the DFT results. In this work, we use coupled-cluster methods with single and double excitations (CCSD) and complete active space self-consistent-field (CASSCF) methods as our computationally expensive methods.

FM 23.4 Mon 17:45 3043

Improving the dynamics of quantum sensors with reinforcement learning — Jonas Schuff, •Lukas Fiderer, and Daniel Braun — Eberhard-Karls-University Tuebingen

Quantum sensors so far have been based almost exclusively on integrable systems, such as precessing spins or harmonic oscillators (e.g., modes of an electro-magnetic field). Non-classical initial states promise large enhancements in measurement precision but are experimentally very difficult to prepare and protect against decoherence.

We recently proposed a new approach that achieves quantum enhancements by rendering the dynamics of the quantum sensor chaotic while using classical initial states that are easy to prepare. Starting from an integrable sensor, the dynamics can be rendered chaotic by applying nonlinear kicks during the parameter-encoding transformation. In this work we deal with the following question: Given the possibility of applying non-linear kicks, what the best strategy to choose the position and strength of these kicks?

This a difficult optimization problem which we tackle with reinforcement learning. As a reward for the learning agent we calculate the quantum Fisher information. At the example of a spin subjected to superradiant damping, we demonstrate how the agent is able to find new strategies. Most strikingly, it is able to adopt to the superradiance decoherence model: quantum Fisher information can be increased further even when it would decay to zero for sensor dynamics without kicks.

FM 23.5 Mon 18:00 3043

Photonic architecture for reinforcement learning — •Fulvio Flamini, Arne Hamann, Sofiène Jerbi, Lea M. Trenkwalder, Hendrik Poulsen Nautrup, and Hans J. Briegel — Institut für Theoretische Physik, Universität Innsbruck, Technikerstraße 25, 6020 Innsbruck, Austria

Artificial intelligence and photonic technologies are driving the limits of present computing devices. Motivated by the recent success in both fields, this work will bring together their state of the art within the framework of reinforcement learning (RL). Specifically, we present the blueprint for a photonic implementation of an active learning agent that can accommodate well-established RL algorithms, such as SARSA, Q-learning, and projective simulation. We numerically investigate its performance within typical task environments, demonstrating that the approach is effective at solving standard RL problems. The simulation is carried out considering imperfect experimental implementations, where we observe that realistic levels of noise can be tolerated or even be beneficial for the learning process. The proposed architecture, based on single-photon evolution on a mesh of tunable beamsplitters, is simple, scalable, and a first integration in portable systems appears to be within the reach of near-term technology.

FM 23.6 Mon 18:15 3043

Machine Learning on Near-Term Universal Quantum Computers — •Manuel Rudolph[1,2], Fred Jendrzejewski[1], and Sebastian Schmitt[2] — [1]Kichhoff-Institute for Physics, Heidelberg, Germany — [2]Honda Research Institute Europe GmbH, Offenbach/Main, Germany

Implementing near-term quantum computers with a small number of qubits and imperfect gate fidelities for real world challenges has been a flourishing field of research in recent years. Quantum-classical hybrid algorithms with shallow quantum circuits for state preparation are being used with success in fields like quantum chemistry and machine learning. This work focuses on the use of near-term quantum computers for unsupervised machine learning on classical data sets with different model infrastructures. It is shown that the quantum state is able to learn the statistics and correlations of data using shallow vari-

ational state preparation. Simple data sets are used to study general aspects such as learning, sampling and generalization of such quantum machine learning implementations in search of practical applications for small quantum machines.