Quantum machine learning: on the design, trainability and noise-robustness of near-term algorithms
Skolik, A.

CHAPTER 8

Conclusion

Throughout this thesis, we have investigated how to design and train variational quantum machine learning models, with a focus on reinforcement learning. We identified four main areas that influence the performance of VQAs in Chapter 1: i) the way in which data is fed in to and read out from the circuit, ii) the structure of the ansatz of the PQC that is used, iii) the classical optimization method used to find the parameters of the circuit, and iv) the influence of noise present on quantum hardware on the training of variational models. We dedicated a chapter to each of these areas.

In Chapter 4, we studied how the classical optimization routine used in a VQA can help to mitigate the barren plateau phenomenon for certain types of circuits. For this, we introduce a method to build the ansatz in an iterative fashion, and partition the set of parameters that are trained in the circuit in a way that prevents the onset of barren plateaus during the optimization routine. We compared this approach to the standard technique of training a fixed ansatz and updating all parameters in every step, and found that our algorithm achieves a lower error on the test set, as well as requiring less wall-clock time considering a realistic sampling rate of a current device.

We address the question of how the data-encoding technique, as well as the choice of observables to read out actions from a quantum Q-learning agent where the PQC is used as a Q-function approximator, influence its performance on two benchmark environments from classical RL literature in Chapter 5. In addition to this, we establish a theoretical separation between classical and quantum learners in a Q-learning setting, and perform an in-depth empirical comparison between quantum Q-learning and the classical DQN algorithm, where we find that the quantum model achieves the same performance as its classical counterpart with a fraction of the parameters.

After this, we address the question of how to design problem-tailored ansatzes for a certain input data type, namely weighted graphs, in Chapter 6. In this chapter, we introduce an ansatz that is equivariant under permutations of the nodes of the input graphs, meaning that the outcome of the PQC does not depend on the order in which representations of nodes are fed into the circuit. We establish a connection to the field of classical geometric deep learning, that is concerned with the design of efficient NN architectures that preserve certain symmetries. We study our ansatz in the context of a learning task on graphs, where a QML model is trained to solve instances of a combinatorial optimization problem by using reinforcement learning. First, we theoretically study the expressivity of our model and find that for our ansatz at depth one, there exists a setting of parameters, for arbitrarily sized instances of the optimization problem that we study, so that our model produces the optimal solution. After establishing that our model is theoretically expressive enough to solve instances of the given optimization problem, we numerically compare our ansatz to general hardware-efficient ansatzes that are unrelated to the problem structure, and find that our equivariant ansatz outperforms them by a large margin.

Finally, we turn to the question of how noise that is present on quantum hardware influences the performance of variational RL models for policy gradients and Q-learning. This study is motivated by a common folklore in the QML community, that conjectures that variational QML models are robust to hardware noise to some degree, due to the classical parameter optimization scheme. In addition, certain results in the classical literature of training neural networks hint at the possibility that a small amount of noise can even be used as a method to combat overfitting. We analytically and numerically study the effect of various noise sources present on quantum hardware: shot noise that is based on the probabilistic nature of quantum measurements, coherent errors due to imperfect control of the quantum device, and incoherent errors that occur due to the device's interaction with its environment. We find that there indeed exists a regime where noise does not prevent quantum RL agents from successfully performing in a given environment, and that there exist cases when training under noisy conditions leads to more robust policies. Additionally, we provide an algorithm to flexibly determine the number of shots required for estimating Q-values, such that the overall number of shots is reduced compared to setting a fixed number of measurements for each circuit evaluation when training a Q-learning agent.

As alluded to in Chapter 1, the goal of this thesis is to contribute to the knowledge of how to successfully train variational QML models, with the aim to foster empirical studies of areas where these types of models can eventually become useful in the future. This is motivated by theoretical results that show that there, indeed, are functions that can only be efficiently learned in a quantum setting, but which are somewhat contrived and not applicable to real-world problems, as well as the historical development of heuristics and machine learning, which showed that often large progress is made when improved hardware becomes available. We hope that, similarly to the development of classical deep learning, the availability of large-scale quantum hardware will lead to the discovery of interesting and valuable applications of quantum machine learning.

> *"We are only one creative algorithm away from valuable near-term applications."*
>
> – Arute et al. [2]