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

Skolik, A.

Summary

Variational quantum machine learning models are often described as the quantum analog of classical neural networks due to the similarity in their training procedure, and are therefore also referred to as quantum neural networks. Unlike for their classical counterparts however, there are still numerous open questions about how to design trainable and performant quantum neural networks. Examples of this include the questions of how to encode classical data into a quantum model, how to structure the gates in the circuits that are used to implement models, and how to avoid pitfalls in the trainability of these models that are unique to the quantum setting. Assuming that similarly to the history of classical machine learning, the development of more performant quantum hardware will facilitate large-scale empirical studies on the usefulness of variational quantum machine learning, it is of key importance to build an understanding of how these models can be trained successfully. This thesis aims to contribute to this understanding by studying various aspects of training variational quantum machine learning models.

We start by giving a basic introduction to the topics of quantum computing, machine learning, and their intersection in Chapters 2 and 3, respectively. In Chapter 4, we study how a fundamental issue in the training of variational quantum circuits, namely barren plateaus in the training landscapes, can be addressed by the classical training algorithm to aid scaling up the size of quantum models. To this end, we provide a training scheme that alleviates the problem of barren plateaus for specific cases and compare it to standard training procedures in the existing literature. While this type of training procedure can in principle be used for arbitrary types of machine learning, we focus our attention on a specific type of learning in subsequent chapters, namely on RL. First, we study in Chapter 5 how the architectural choices made for a PQC-based quantum agent influence its performance on two classical benchmark tasks from RL literature, where we specifically consider the question of encoding data into, and reading information out of the quantum model. In addition, we establish a theoretical separation between classical and quantum models for the specific type of RL algorithm that we use, and also perform an in-depth empirical comparison of the quantum model developed in our work to a classical neural network that performs the same task. In addition to the questions of how to encode data and read out information from a PQC, the third key question in the performance of a variational quantum machine learning model is how to design the structure of the circuit itself, also referred to as the ansatz. For this reason, we move on to study this question in Chapter 6 and introduce an ansatz that is tailored to a specific type of input data, namely to weighted graphs. To do this, we take inspiration from the classical field of geometric deep learning, and design a PQC that preserves an important symmetry in graph-based input data. We analytically study the expressivity of this type of circuit, and then go on to numerically compare it to ansatzes that are not tailored to the specific training data at hand. Finally, another important consideration in the study of algorithms for the NISQ era is how the given learning algorithms and models are influenced by quantum hardware-induced noise. In Chapter 7, we study this for two of the variational RL paradigms from recent literature. We investigate analytically and numerically how various types of errors, namely coherent, incoherent, and measurement-based errors, affect the training performance of variational RL algorithms and the robustness of the learned policies. In particular, this study includes an evaluation of the performance of the models we introduced in Chapter 5 and Chapter 6 under various types of noise that are expected to be present on near-term hardware.

With the above, this thesis aims to contribute to building a foundation of knowledge about how to successfully train variational quantum machine learning models, in the hope that similarly to classical machine learning, this knowledge will one day, when quantum hardware has sufficiently matured, aid demonstrations of the practical usefulness of these types of algorithms.