Quantum methods for machine learning and classical dynamics

Barthe, A.M.

CHAPTER 8

Conclusion

We summarize our answer to the first research question:

Research Question 1
What are the minimal resources needed to achieve universality with parameterized quantum circuits?

In Chapter 3, we have refined the understanding of the expressivity of Quantum re-uploading models relative to the number of reuploads L. Precisely, we have shown that quantum reuploading models with repeated layers, follow a Gaussian profile with a width proportional to $\sqrt{L}$. This means that while having in theory a frequency support growing linearly with L, in practice it grows quadratically more slowly. Whether this is a bug or a feature depends on the use case: on the one hand, it enhances generalization performance through Lipschitz regularization, on the other hand, it cannot capture function with fine-details.

In Chapter 4, for continuous distribution modeling, we proved that a type of quantum generative model, that we call *expectation-value samplers* can achieve full universality for continuous multivariate distributions. We gave resource bounds in terms of qubits and number of shots, derived from information-theoretic arguments (e.g., variations of the Holevo bound). We showed that such bounds are tight through the existence of a family of circuits that saturate them. Our results highlight a clear trade-off: more

8

qubits reduce measurement cost and vice versa, but minimal resources must scale polynomially in the problem dimension.

We summarize our answer to the second research question:

Research Question 2
How can we leverage our knowledge of parameterized quantum circuits to construct practical provable quantum learning advantages?

In Chapter 5, we formalized a supervised learning task of learning PQC-based functions when gates with unknown parameters are present in the circuit. We showed that this task can yield provable exponential classical-quantum learning advantages under common complexity assumptions. At the core of our proposed quantum learning algorithm is a new subroutine we call *Fourier Coefficient Extraction*, inspired by the Fourier analysis of PQCs. Using trotterization, we connected this task to the more physically relevant task of learning Hamiltonian dynamics, where the input is a classical data set of input-output pairs from time evolutions with unknown parameters. We note that, using the same process, any BQP-complete problem that can be compiled as a PQC can be used to construct a learning task with a learning separation.

We summarize our answer to the third research question:

Research Question 3
How does the computational complexity of bosonic systems relate to that of classical systems and qubit-based systems?

The algorithm we propose in Chapter 6 to solve polynomial differential equations yields a compilation of primitive bosonic gates to approximate the evolution according to the system dynamics. We note that this compilation is exponential in the degree of the polynomial, indicating that non-polynomial equations (such as gravitational dynamics) yield harder simulations. The level of squeezing of the initial bosonic state controls the spread around the initial condition, therefore the more photons are injected in the system, the more precise the simulation is. This uncertainty is particularly critical when simulating chaotic systems.

In Chapter 7, we focus on Gaussian states on exponentially many modes. We define a decision problem that is shown to be BQP-complete when considering only a small subset of interferometers. We prove that squeezing gates are equivalent to imaginary time evolution, and therefore, their addition boosts the complexity to PostBQP-complete.

All provable advantages demonstrated in this thesis, whether for computational or learning tasks, require deep, reliable circuits with high connectivity that are only feasible on fault-tolerant quantum computers. Current Noisy Intermediate-Scale Quantum (NISQ) devices lack the error correction and coherence necessary to execute such algorithms reliably. While they may support proof-of-concept implementations and limited-scale testing, they are insufficient for realizing scalable, useful, and provable quantum advantage. The results presented in this thesis instead characterize the minimal resources and algorithmic structures necessary for achieving such advantages, delineating a clear target for the capabilities of future quantum hardware.