Quantum methods for machine learning and classical dynamics

Barthe, A.M.

CHAPTER 1

Introduction

1.1. Preface

Computation is a critical tool for modern society, enabling everything from scientific discovery to secure communication and artificial intelligence. As our ambitions grow, so do the demands on our *computational resources*, and some problems remain stubbornly out of reach for even the most powerful classical computers. Simultaneously, we have entered a new paradigm where computation is increasingly driven by data feeding into machine learning methods. Today's *machine learning* systems are ubiquitous, from image recognition to language models, and rely on training over vast datasets. This change created new algorithmic challenges and placed intense pressure on hardware and software alike to scale effectively.

Quantum computers have emerged as an avenue to address computational problems that are hard for classical computers by leveraging the principles of *quantum mechanics*. The key to quantum advantage is the ability to exploit quantum phenomena like interference and entanglement. Quantum algorithms can be designed to amplify the probability of correct answers and suppress the wrong ones, manipulating the quantum state as a vector in a high-dimensional space. The operations used to manipulate quantum states are called *quantum gates*, and they can be combined to form *quantum circuits*. Quantum circuits are the quantum analog of classical circuits, and they can be used to perform some computations faster than classical computers. For some problems, such as unstructured search, quantum computers lead to a quadratic speed-up [1]; for others, like factoring [2], the speed-up can be exponential. But quantum computers are not a universal solution to address all computational problems efficiently. As we explore the potential of quantum computing and look for *quantum advantage*, it is crucial to maintain a clear-eyed view of both its strengths and its limitations.

As quantum advantages have been theoretically proven for specific computational problems, the question was extended to whether advantages could also be achieved in *learning tasks* [3]. One attempt at machine learning making use of quantum computers took the form of *variational algorithms* [4]. These are hybrid methods where a quantum circuit, composed of gates governed by a set of adjustable parameters, is used in tandem with a classical optimizer to minimize a cost function, often representative of the prediction error. Unlike traditional quantum algorithms, like Shor's [2] or Grover's [1], which offer provable guarantees under clear assumptions, most variational algorithms are heuristic by nature. They typically come with no formal guarantees about their convergence, accuracy, or advantage over classical methods, but are designed to perform well in practice. In fact, one motivation for variational algorithms was to work within the severe constraints of early quantum hardware [5], where obtaining reliable results is only possible for short circuits.

At the heart of quantum variational algorithms are *parameterized quantum circuits* (PQCs): quantum circuits whose structure is fixed, but which include tunable quantum gates. These gates typically depend on continuous parameters, like rotation angles, that can be optimized to steer the quantum state toward a useful configuration. The idea is simple enough: start with a tunable quantum model, tweak the parameters using feedback from measurements, and hope that this loop uncovers a useful solution. Whether and when this actually leads to a quantum advantage is still an open question, but the framework itself became a foundation for much of today's quantum machine learning research [6–9]. In what follows, we will discuss the challenges and limitations of PQCs in more detail.

Most variational algorithms rely on estimating *expectation values*, real numbers that represent the average outcome of a measurement on a quantum state. In practice, these are not computed directly but estimated statistically by repeatedly measuring the output of a quantum circuit, each time producing a random bitstring. The accuracy of this estimate improves with the number of repetitions (called shots), but only slowly: the standard error decreases as $1/\sqrt{N}$, where N is the number of shots.

While quantum techniques like amplitude amplification [10] can sometimes offer quadratic improvements, the output of a quantum circuit remains fundamentally random. In contrast, classical computing typically allows for deterministic evaluation, and higher accuracies may be reached more easily. Adding a bit in a floating-point number can double the precision of most calculations (except, for example, in Markov chain Monte Carlo methods). This fundamental difference means that achieving high accuracy in variational algorithms often comes at a significant sampling cost.

Another major challenge with PQCs is *trainability*, that is, the ability to find parameter values that solve the problem at hand. The loss function measures how well the quantum model performs on a given task and is typically minimized using gradient-based methods like gradient descent. These methods require computing the derivative of the loss with respect to each parameter, which tells the algorithm in which direction to update the parameters to improve performance. However, many PQC architectures suffer from the so-called *barren plateau* problem [11], where the gradients vanish exponentially with the number of qubits or circuit depth. In these cases, the signal becomes indistinguishable from shot noise, and optimization cannot proceed. As a result, the search for architectures that provably avoid barren plateaus has become a central focus in the design of trainable quantum models.

In parallel, advances in classical simulation techniques had direct consequences for quantum machine learning. Significant progress was made in *classical simulation* of quantum circuits, that is, computing their outputs using classical algorithms without needing access to a quantum device. In general, this is a hard problem: the resources needed for a classical computer to naively track the full quantum state grow exponentially with the number of qubits. However, some quantum circuits are easy to simulate. This is often the case when the circuit is shallow (i.e., has few gates), or when it is restricted to specific gate sets, such as the Clifford group [12]. Efficient simulation is also possible when circuits have special algebraic structure [13], or when they are highly noisy [14–16], making either their average output easy to simulate or their output distribution effectively easy to sample from. Understanding which circuits are classically simulable is essential because these circuits cannot provide a quantum advantage by nature.

In particular, many of the PQC architectures that were specifically designed to avoid barren plateaus were later shown to be classically simulable [13, 17]. This exposed a growing tension: circuits that are easy to train are often also easy to simulate classically, and therefore unlikely to provide any true quantum advantage [18]. This narrowing gap between trainability and simulability has forced the community to critically reassess the role of PQCs in quantum machine learning. While the early enthusiasm has been tempered by these findings, the search for quantum learning advantages continues, looking for learning algorithms that not merely *use* quantum computers but rather *need* quantum computers.

Despite these setbacks, there is rigorous evidence that quantum computers can outperform classical ones on certain learning tasks. These results are known as *learning separations*: formal examples where a quantum learner can efficiently learn a function, while any classical learner would require super-polynomially more resources. The first such separations were based on cryptography-inspired tasks such as the factoring problem for Blum integers [19] and the discrete logarithm problem [20]. The latter is a central problem in cryptography and is at the heart of most protocols securing communication today. More recent works have uncovered learning separations that go beyond cryptographic tasks [21, 22], showing that quantum advantage can emerge when the target function relates to the set of hardest problems quantum computers are expected to solve efficiently. These are valuable because they show that a quantum advantage in learning is possible in principle, but this comes with a caveat: how these problems may relate to physically relevant tasks remains an open question. It is a challenge to find quantum learning separations that have an impact on machine learning applications we care about for real-world datasets.

This thesis has the overarching goal of looking for quantum advantages. While this theme is broad, to contribute to its elucidation we will focus on two main aspects. On the one hand, deepening our understanding of the capacities of quantum machine learning models, exploring the universality of PQC-based models, and finding provable quantum learning advantages. And on the other hand, defining and analyzing the computational complexity of simulating classical and quantum systems with bosonic systems. We detail these two aspects in the following section, which also outlines the research questions that guide this thesis.

1.2. Research questions

In order to make the research questions more precise, we briefly clarify some basic terminology, which we will make fully precise later in Chapter 2.

First, we unpack the theoretical capabilities of parameterized quantum circuits (PQCs) as models for two core tasks in learning: supervised prediction and generative sampling. In *supervised learning*, the model is trained on labeled example pairs to learn an underlying mapping. The learning algorithm should provide a model that predicts outputs close to the correct ones for unseen inputs. By contrast, *generative modeling* is given unlabeled samples and aims to learn the distribution of the data itself. The learning algorithm should provide a model that generates new samples from a distribution that is as close as possible to the target distribution.

Expressivity is a fundamental concept across both supervised and generative learning: it quantifies a model's ability to represent all functions or probability distributions. Formally, given a target space and a model class, we say a model is more expressive if it can approximate a larger portion of the target space. The highest level of expressivity is called *universality*, where the model class is dense in the target space: any target element can be approximated up to arbitrary precision by some model element. This property is well-known in classical settings. For example, neural networks satisfy universal approximation theorems under mild conditions [23, 24].

This opens the question of whether PQCs can achieve universality, and if so, under what conditions. Some universality results have been established for PQCs, showing that they can approximate any continuous function [25, 26]. In the context of both supervised learning and generative learning, we sharpen this question in the case of limited resources, and this leads us to the first research question.

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

Secondly, we turn our attention to the question of when quantum computers can provide a provable learning advantage over classical ones. As mentioned in the preface, most variational methods are heuristic and lack theoretical guarantees. Yet, there are notable exceptions where quantum learning advantage can be proven rigorously. These proofs rely on statistical learning theory, specifically the *probably approximately correct* (PAC) learning framework [27, 28]. The PAC framework provides a formal definition for a learning algorithm to be successful, by upper-bounding the probability that the algorithm exceeds an error threshold on unseen data. If the resources in terms of computation time and size of dataset needed for the algorithm to be successful are polynomial in the size of the training set, we say that the class is efficiently PAC-learnable. This framework gives us a rigorous benchmark to claim that a quantum learning algorithm is efficient, moving us beyond heuristics.

If, in addition to being efficient, the quantum learning algorithm can learn a function that is provably hard for any classical algorithm to learn, we say that we have a *learning separation*. As mentioned in the preface, several learning separations were proven, first on cryptographic tasks [19, 20] and later for tasks associated with PromiseBQP-complete problems [21, 22]. While these previous works did not focus on PQCs, in this thesis, we investigate PQC-based methods that allow a learning separation, leading us 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?

So far, we have focused our attention on the advantages provided by qubit-based quantum computers; but *continuous-variable quantum computing* (CVQC) constitute an alternative computational model. CVQC leverages quantum systems with continuous degrees of freedom, such as the position and momentum of particles, to perform computation. Alternatively, such states can be understood as vectors in an infinite-dimensional space, called the Fock space, which is in contrast to the discrete qubit-based model, where quantum information is encoded in two-level systems (qubits). This paradigm typically arises from the physics of bosonic systems, where quantum information is encoded in the states of harmonic oscillators, such as modes of light or vibrational modes in ions. Within CVQC, a prominent role is played by *Gaussian states*, whose quadrature operators measurements can be described as Gaussian distributions. Operations called Gaussian gates (e.g., displacements, rotations, and squeezers) preserve the Gaussian property of states. Gaussian states evolved under Gaussian gates on n modes are classically simulable. However, adding non-Gaussian gates lifts this restriction and enables universal quantum computation in the continuous-variable regime [29]. The complexity of unrestricted bosonic systems has been studied in [30] and related to complexity classes such as EXP. In the final part of this thesis, we turn to bosonic systems with some added restrictions and relate their computational complexity to that of classical systems and qubit-based quantum computers

In Chapter 2, we introduce the complexity classes mentioned in this paragraph at greater length, but we provide a high-level introduction here first. We briefly mentioned in the preface the set of hardest problems that quantum computers are expected to solve efficiently. These are known as *PromiseBQP-complete* problems, that is the hardest problems that are in BQP (Bounded-error Quantum Polynomial time). BQP [31] is the complexity class of all decision problems that can be decided by a quantum computer in polynomial time with bounded error. If any of the BQP-complete problems was to be solved efficiently on a classical computer, this would violate widely believed complexity theory assumptions. Beyond BQP, there are even harder quantum complexity classes, such as QMA (Quantum Merlin-Arthur), which is the quantum analog of NP [32, 33]. More powerful still is PostBQP, which allows post-selection on specific measurement outcomes—conditioning on events with exponentially small probability [34]. Remarkably, PostBQP includes QMA and coincides with the classical complexity class PP [34], situating it as a powerful tool to probe the limits of quantum computation.

This leads us to the following research question, which explores how the complexity of bosonic systems maps onto the complexity landscape of classical and qubit-based systems.

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

1

1.3. List of contributions

The following works form the basis of this dissertation.

[35] A. Barthe and A. Pérez-Salinas. Gradients and frequency profiles of quantum re-uploading models. Quantum, vol. 8, p. 1523 (2024).

[36] A. Barthe, M. Grossi, S. Vallecorsa, J. Tura, and V. Dunjko. Parameterized quantum circuits as universal generative models for continuous multivariate distributions. npj Quantum Inf, vol. 11, no. 1, p. 121 (2025)

[37] A. Barthe, M. yaghubi, M. Grossi, and V. Dunjko. Quantum Advantage in Learning Quantum Dynamics. arXiv:2506.17089 (2025).

[38] A. Barthe, M. Grossi, J. Tura, and V. Dunjko. Continuous variables quantum algorithm for solving ordinary differential equations. IEEE International Conference on Quantum Computing and Engineering (2023)

[39] A. Barthe, M. Cerezo, A. T. Sornborger, M. Larocca, and D. García-Martín. Gate-Based Quantum Simulation of Gaussian Bosonic Circuits on Exponentially Many Modes. Phys. Rev. Lett., vol. 134, no. 7, p. 070604 (2025).

The following publications were co-authored during the course of the PhD and are not included in this thesis.

[40] M. Sahebi, A. Barthe, Y. Suzuki, Z. Holmes, and M. Grossi. On Dequantization of Supervised Quantum Machine Learning via Random Fourier Features. arXiv: arXiv:2505.15902 (2025).

[41] A. Pérez-Salinas, M. Y. Rad, A. Barthe, and V. Dunjko. Universal approximation of continuous functions with minimal quantum circuits. arXiv: arXiv:2411.19152 (2025).

1.4. Overview of this thesis

Chapter 3 is based on the peer-reviewed journal article [35], the first on the list in the previous section. In this chapter, we analyze the average behavior of a class of PQCs called quantum re-uploading models. These models iteratively encode data into quantum circuits, with trainable gates in between. We focus on theoretical aspects of QRU focusing on their trainability and their expressivity in a supervised learning context. We study the gradients of cost functions associated with these models and quantify the influence of data encoding layers on their trainability. Furthermore, we explore the expressivity of quantum re-uploading models by examining properties of the Fourier decomposition of the models. Our results demonstrate that, on average, this decomposition follows a Gaussian profile and that high-frequency components vanish.

Chapter 4 is based on the peer-reviewed journal article [36], the second on the list in the previous section. In this chapter, we move to a generative modeling context for multidimensional distributions of continuous random variables. We formalize a class of quantum generative models for such distributions, which we call *expectation value samplers*. The idea is to use a PQC to sample from a distribution defined by the expectation values of a set of observables, which can be efficiently computed using quantum circuits. In this chapter, we derive tight necessary and sufficient conditions on the quantum resources required to achieve universality for such models. We do so by leveraging existing knowledge on parameterized quantum circuits, demonstrating how these models can represent complex distributions. Furthermore, we identify a physically relevant distribution that these models can efficiently represent.

Chapter 5 is based on the preprint [37], the third on the list in the previous section. Building on the refined understanding of PQCs and their connection to Fourier analysis, we investigate how quantum feature maps can be constructed to achieve provable learning separations. Feature maps are a key component of many machine learning algorithms, transforming the input data into a higher-dimensional space where it is easier to learn patterns. In this chapter, we introduce an algorithm for extracting the Fourier spectrum of PQC-based functions, enabling the design of feature maps that are likely intractable for classical computers. Leveraging this approach, we demonstrate a quantum learning advantage for a specific learning task based on PQC-based functions. Furthermore, we extend this framework to derive a learning separation for a physically relevant task involving Hamiltonian time evolution. These results highlight the potential of PQCs in constructing quantum learning advantages for tasks grounded in quantum dynamics, while emphasizing the necessity of fault-tolerant quantum computation for such applications.

Chapter 6 is based on the peer-reviewed proceedings paper [38], the fourth on the list in the previous section. We propose a method for using continuous variable quantum computers to solve nonlinear differential equations, leveraging the Koopman von Neumann framework. By working directly in the continuous variable quantum computing paradigm, we avoid the truncation errors associated with finite-dimensional approximations. Our algorithm compiles a sequence of Gaussian and non-Gaussian gates to approximate the time evolution of the Koopman von Neumann Hamiltonian for polynomial ordinary differential equations. This method is particularly suited for solving initial distribution problems, where the evolution of probability distributions is studied rather than individual initial conditions.

Chapter 7 is based on the peer-reviewed journal article [39], the fifth on the list in the previous section, with additional unpublished results on larger complexity classes. Quantum computers are known to provide advantages for simulating certain large classical dynamics. Inspired by quantum algorithms for exponentially many coupled oscillators [42], we develop a quantum simulation framework for bosonic Gaussian states on exponentially many modes. This framework encodes the first and second moments of quadrature operators into qubit states, enabling efficient simulation of Gaussian bosonic circuits using gate-based quantum computers. We introduce a mapping between bosonic gates and qubit gates, allowing particle-preserving Gaussian bosonic circuits to be simulated efficiently. Furthermore, we identify a PromiseBQP-complete problem within this framework, demonstrating that simulating Gaussian bosonic circuits on exponentially many modes is as powerful as universal quantum computation. Finally, we show that adding a layer of squeezing gates to interferometers boosts the complexity of the problem to PostBQP-complete.