



Universiteit
Leiden

The Netherlands

Computational speedups and learning separations in quantum machine learning

Gyurik, C.

Citation

Gyurik, C. (2024, April 4). *Computational speedups and learning separations in quantum machine learning*. Retrieved from <https://hdl.handle.net/1887/3731364>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3731364>

Note: To cite this publication please use the final published version (if applicable).

Chapter 1

Introduction

In this chapter, we provide an introduction to the topics discussed in this thesis. We start by discussing the problem statement and research questions covered by this thesis, which can be found in Section 1.1. Afterwards, in Section 1.2, we give a detailed overview of the contributions of this thesis.

1.1 Problem statement and research questions

Quantum machine learning (QML) is a rapidly growing field that has brought forth numerous proposals regarding ways for quantum computers to help analyze data. Several of these proposals involve using quantum algorithms for linear algebra – most notably Harrow, Hassidim and Lloyd’s matrix inversion algorithm [100] – to exponentially speed up tasks in machine learning. Other proposals involve using parameterized quantum circuits [102, 167, 31] to identify novel quantum learning models that are better suited to the limitations of near-term quantum computing, rather focussing on improving established classical methods. These QML proposals have all been hailed as possible examples of quantum computing’s “killer application”: genuinely and broadly useful quantum algorithms that (superpolynomially) outperform their best known classical counterparts. In this thesis, we study several of these QML proposals and we investigate whether and how they are able to provide (superpolynomial) speedups over their classical counterparts, and how to get the best possible performance out of these proposals. Specifically, the problem statement of this thesis is:

Problem statement. *Can we provide evidence that various QML proposals can (superpolynomially) outperform their classical counterparts, and what methods can we devise to attain their best possible performance?*

The first QML proposal we study concerns quantum algorithms for topological data analysis, as first introduced by Lloyd, Garnerone and Zanardi [130]. Specifically, in Chapter 3, we study the potential for quantum algorithms to achieve superpolynomial speedups for linear-algebraic problems in topological data analysis. An important linear-algebraic problem that arises in topological data analysis is that of computing

so-called *Betti numbers*, which can be formulated as computing the dimension of the kernel of a set of matrices whose entries encode the connectivity of cliques (i.e., complete subgraphs) in a given graph. While Lloyd, Garnerone and Zanardi developed a quantum algorithm that in certain regimes is able to solve this problem superpolynomially faster than the best-known classical algorithm [130], it is unclear whether this speedup will persist with the development of new classical algorithms. In particular, for several other linear-algebraic QML proposals the previously speculated superpolynomial speedups were revealed to actually be at most polynomial speedups, as exponentially faster classical algorithms were devised that operate under analogous assumptions [186, 60] (an event called “dequantizations”). While polynomial speedups have appeal on paper, an analysis involving near-term device properties revealed that low-degree polynomial improvements are not expected to translate to real-world advantages due to various overheads [27]. Thus, finding superpolynomial speedups is of great importance, especially in the early days of practical quantum computing. Therefore, we examine the linear-algebraic QML algorithms for Betti numbers and study whether speculated superpolynomial quantum speedups will not be lost due to development of better classical algorithms.

Research question 1. *Can the linear-algebraic QML algorithms for Betti numbers maintain their speculated superpolynomial quantum speedups, even with the development of better classical algorithms?*

Next, we turn our focus to QML proposals that involve the use of parameterized quantum circuits [102, 167, 31] to build genuinely new quantum machine learning models. Specifically, in Chapter 4, we study how to optimally tune a family of these new quantum machine learning models based on the principles of *structural risk minimization*. In the context of structural risk minimization, it is crucial to identify the tunable aspects of a model, i.e., hyperparameters, that impact both its performance on training data and its generalization performance. An important part of this investigation involves characterizing complexity measures, such as the VC-dimension or fat-shattering dimension, associated with these models and understanding how hyperparameters influence these complexity measures. Consequently, we focus on characterizing complexity measures of novel quantum learning models based in parameterized quantum circuits. In particular, we aim to identify the hyperparameters that exert influence over these complexity measures, a critical step in effectively implementing structural risk minimization.

Research question 2. *Can we identify hyperparameters within novel quantum learning models based on parameterized quantum circuits that impact both complexity measures and performance on training data, as is crucial for the successful implementation of structural risk minimization?*

Afterwards, in Chapter 5, we study how these new quantum machine learning models based on parameterized quantum circuits [102, 167, 31] can be used in the field of reinforcement learning. Reinforcement learning is a flavour of learning where one has to learn through interacting with an environment and adjusting its behaviour based on rewards obtained (i.e., there is no large amount of labeled data available). Arguably, the largest impact quantum computing can have is by providing enhancements to the

hardest learning problems. From this perspective, reinforcement learning stands out as a field that can greatly benefit from a powerful hypothesis family. Nonetheless, the true potential of near-term quantum approaches in reinforcement learning remains very little explored. The few existing works [58, 131, 202, 110] have failed so far at solving classical benchmarking tasks using PQCs and left open the question of their ability to provide a learning advantage. Consequently, we focus on designing quantum machine learning models based on parameterized quantum circuits that are on par with the best classical models in standard benchmarking tasks, but also outperform all classical models in certain other tasks.

Research question 3. *How can new quantum machine learning models based on parameterized quantum circuits be effectively leveraged within the realm of reinforcement learning? Specifically, can these quantum approaches demonstrate the potential to be on par with classical models in standard benchmarking tasks and outperform them in novel specific scenarios?*

Finally, in Chapter 6, we tackle the challenge of identifying learning problems that exhibit provable exponential speedup for quantum learning algorithms compared to classical learning algorithms. The first thing we address is that there is no single definition of what precisely constitutes a *learning* separation. In particular, when trying to come up with a definition there are many choices to be made, and various choices make sense depending on the particular settings. This ambiguity can lead to conflating the task of learning in an intuitive sense with a purely computational task. Moreover, we study existing learning separations [126, 173] and carefully delineate where the classical hardness of learning lies and the types of learning separations they achieve. Next, we set out to find new examples of learning separations where the classical hardness lies more in learning in an intuitive sense rather than evaluating the functions to be learned. Finally, we turn our attention to the folklore in the community that states that quantum machine learning is most likely to have advantages when the data is quantum-generated. However, it is not immediately clear how quantum-generated data can give rise to learning separations. We address this question by exploring the additional complexity-theoretic assumptions required to build such a learning separation.

Research question 4. *How can we identify learning problems that exhibit a provable exponential speedup for quantum learning algorithms compared to their classical counterparts, and can we confirm the validity of the folklore that quantum machine learning excels when handling quantum-generated data?*

1.2 List of contributions

- **Chapter 3:** We show that the linear-algebraic methods underlying the algorithm of Lloyd et al. are “safe” against general dequantization approaches of the type introduced in [186, 60], and that the corresponding computational problem (i.e., a generalization of Betti numbers) is classically intractable under widely-believed assumptions. Specifically, we show that a natural generalization of the Betti number problem is hard for the complexity class DQC1 (see Section 3.2).

Since the hard problems in DQC1 are widely-believed to require superpolynomial time on a classical computer, this ensures that our generalization also most likely requires superpolynomial time on a classical computer. This further establishes the potential of these methods to be a source of useful quantum algorithms with superpolynomial speedups over classical methods. We concretely demonstrate the potential of these methods by connecting them to practical problems in machine learning and complex network analysis (see Section 3.3). Finally, we provide examples of instances that satisfy the requirements for the quantum algorithm to be efficient, while making sure the best-known classical algorithms are not efficient (see Section 3.2.4).

- **Chapter 4:** By exploiting a connection between the new quantum learning models and certain established classical learning models (i.e., linear-classifiers) we characterize some complexity measures of the new quantum learning models (i.e., their VC- and fat-shattering dimension). Since these complexity measures characterize the generalization performance of the machine learning model, our results give rise to ways to fine-tune your model such that you perform well on training data, while ensuring that the generalization performance remains good enough (i.e., the principle of structural risk minimization).
- **Chapter 5:** We exhibit how to use quantum machine learning models based on parameterized quantum circuits to solve problems in reinforcement learning using the policy-gradient algorithm. Next, we build reinforcement learning settings where we (i) can prove that the quantum learning model performs much better than any classical learner, and (ii) can provide numerical evidence that the quantum learning model performs much better than any deep neural network based learner (i.e., the current classical state of the art).
- **Chapter 6:** We delve into the nuances of computational learning theory and highlight how subtle variations in definitions lead to distinct requirements and tasks for learners. Next, we examine existing learning problems demonstrating provable quantum speedups [126, 173] and observe their reliance on the classical complexity of evaluating the data-generating function rather than its identification. To address this limitation, we present two novel learning scenarios where the primary classical challenge lies in identifying the underlying function generating the data. Additionally, we explore computational hardness assumptions that can be utilized to establish quantum speedups in situations where the data is quantum-generated. This implies quantum advantages in various natural settings such as condensed matter and high-energy physics.

This thesis is based on the following papers:

- [1] Casper Gyurik, Chris Cade, and Vedran Dunjko. Towards quantum advantage via topological data analysis. *Quantum*, 6, 2022.
- [2] Dominic W Berry, Yuan Su, Casper Gyurik, Robbie King, Joao Basso, Alexander Del Toro Barba, Abhishek Rajput, Nathan Wiebe, Vedran Dunjko, and Ryan Babbush. Analyzing prospects for quantum advantage in topological data analysis. *PRX Quantum*, 5, 2024.
- [3] Casper Gyurik, Dyon van Vreumingen, and Vedran Dunjko. Structural risk minimization for quantum linear classifiers. *Quantum*, 7, 2023.
- [4] Casper Gyurik and Vedran Dunjko. Exponential separations between classical and quantum learners. arXiv:2306.16028, 2023.
- [5] Sofiene Jerbi, Casper Gyurik, Simon Marshall, Hans Briegel, and Vedran Dunjko. Parametrized quantum policies for reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

In the course of their PhD, the author has additionally co-authored the following articles that are not included in this thesis:

- [6] Simon Marshall, Casper Gyurik, and Vedran Dunjko. High dimensional quantum machine learning with small quantum computers. *Quantum*, 7, 2023.
- [7] Sofiene Jerbi, Casper Gyurik, Simon Marshall, Riccardo Molteni, and Vedran Dunjko. Shadows of quantum machine learning. arXiv:2306.00061, 2023. *Submitted to Nature Physics*.