



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 7

Conclusion

In this chapter, we present the conclusions of the thesis. In Section 7.1, we recall the problem statement and research questions and we provide them with answers. Next, in Section 7.2, we discuss the limitations of our work. Finally, in Section 7.3, we discuss some promising directions for future work.

7.1 Research overview

In this thesis, we studied the speedups provided by several QML proposals over their classical counterparts, and how to get the best possible performance out of these proposals. Specifically, the problem statement of this thesis was:

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

The above problem statement was split into 4 research question. Below, we restate these research questions and subsequently provide them with answers.

Chapter 3 investigated the potential of a class of problems arising from the quantum algorithm for topological data analysis [130] to become genuinely useful applications of quantum computers with a superpolynomial quantum speedup, with the first research question in mind:

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?*

We showed that this algorithm along with a number of new algorithms provided by us (with applications in numerical linear algebra, machine learning and complex network analysis) solve problems that are classically intractable under widely-believed complexity-theoretic assumptions by showing that they are as hard as simulating the one clean qubit model. Specifically, our results showed that the methods of the quantum algorithm for topological data analysis withstand the sweeping dequantization

results of Tang et al. [186, 60]. To analyze whether it is possible to further strengthen the argument for quantum advantage (or, to actually find an efficient classical algorithm) for the narrow TDA problem, we investigated state-of-the-art classical algorithms and we highlighted the theoretical hurdles that, at least currently, stymie such classical approaches. Regarding near-term implementations, we identified that implementing sparse access to the input matrix is a major bottleneck in terms of the required number of qubits, we proposed multiple methods to circumvent this bottleneck via classical precompilation strategies, and we investigated the required resources to challenge the best known classical methods.

Chapter 4 studied how to implement structural risk minimization for quantum machine learning models based on parameterized quantum circuits, to answer the second research question:

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?*

In particular, in Theorem 12 and Theorem 14 we characterized the VC-dimension and fat-shattering dimension of these quantum models and identified hyperparameters – such as the rank and Frobenius norms of the observables – that influence these complexity measures. Moreover, in Proposition 15 and Proposition 17 we showed that these hyperparameters also influence the performance that these quantum models can have on training data. Finally, we showed how our findings can be used to construct new quantum machine learning models with favourable performance guarantees based on the principle of structural risk minimization.

Chapter 5 investigated quantum reinforcement learning agents based on parameterized quantum circuits motivated by our third research question:

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?*

We proposed several constructions and showed the impact of certain design choices on learning performance. In particular, we introduced the SOFTMAX-PQC model, where a softmax policy is computed from expectation values of a parameterized quantum circuit with both trainable observable weights and input scaling parameters. These added features to standard parameterized quantum circuits used in machine learning (e.g., as quantum classifiers) enhance both the expressivity and flexibility of parameterized quantum circuit policies, which allows them to achieve a learning performance on benchmarking environments comparable to that of standard deep neural networks. We additionally demonstrated the existence of task environments, constructed out of parameterized quantum circuit, that are very natural for parameterized quantum circuit agents, but on which deep neural network agents have a poor

performance. To strengthen this result, we constructed several reinforcement learning environments, each with a different degree of degeneracy (i.e., closeness to a supervised learning task), where we showed a rigorous separation between a class of parameterized quantum circuit agents and any classical learner, based on the widely-believed classical hardness of the discrete logarithm problem. We believe that our results constitute strides toward a practical quantum advantage in reinforcement learning using (near-term) quantum devices.

Chapter 6 focused on the identification of learning problems within the probably approximately correct (PAC) learning framework, where quantum learners exhibit exponential advantages over classical learners, as motivated by our fourth research question:

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?*

Firstly, we delved into the intricacies of precisely defining what it means for a quantum learner to exhibit an exponential advantage over its classical counterpart. Subsequently, we studied prior instances of learning separations [126, 173], pinpointing the exact source of classical hardness and the quantum edge. Lastly, we examined the folklore that quantum machine learning excels most in scenarios involving quantum-generated data. In doing so, we established a framework through which any BQP-complete problem can lead to a learning separation, thereby substantiating the quantum advantage across numerous domains in physics.

7.2 Limitations

In this section, we highlight certain limitations in the outcomes of this thesis. Firstly, concerning our findings in Chapter 3, it is important to note that our discussion of the noise-robustness of the quantum algorithm for topological data analysis lacks experiments to confirm or reject our statements due to our limited access to sufficiently large quantum hardware. Secondly, in relation to our results in Chapter 4, our analysis does not consider the structure of the feature map, which has the potential to enhance the effectiveness of structural risk minimization. Moreover, in reference to our outcomes in Chapter 5, our ability to benchmark reinforcement learning models was limited to toy problems, as we lacked access to hardware capable of handling real-world problem sizes. Lastly, with respect to our findings in Chapter 6, it is worth noting that the results of Theorem 26, in most instances, rely on contrived data distributions that may not be representative of real-world scenarios.

7.3 Future work

The findings presented in this thesis open up exciting avenues for future research, offering several promising directions to explore. In this section, we highlight a few of

these potential interesting opportunities for future work.

Regarding the topological data analysis results in Chapter 3, there are several interesting open questions. First, it remains open whether ABNE, as outlined in Section 1.1, is indeed DQC1-hard. In particular, it remains open whether LLSD retains its DQC1-hardness when restricted to combinatorial Laplacians. Exploring quantum algorithms' capabilities for problems in topological data analysis beyond computing Betti numbers, such as handling other aspects of barcodes [83], presents another avenue of research. Lastly, the potential utility of quantum algorithms in computing eigenvalues and eigenvectors of combinatorial Laplacians for complex network analysis, as hinted in [134], merits further investigation.

Concerning the results presented in Chapter 4, which delve into structural risk minimization for quantum machine learning models relying on parameterized quantum circuits, several interesting research directions emerge. Firstly, it is worth exploring alternative complexity measures beyond VC-dimension and fat-shattering dimension. Such an exploration can potentially uncover additional sets of hyperparameters relevant to the structural risk minimization tradeoff. Additionally, there is room for investigating how the phenomenon of overparameterization, as extensively studied in the context of neural networks [18], extends to quantum machine learning models that employ parameterized quantum circuits. This investigation can provide valuable insights into the generalization performance of these quantum models, shedding light on their behavior in comparison to classical counterparts.

In light of the results discussed in Chapter 5, which pertain to reinforcement learning with quantum machine learning models based on parameterized quantum circuits, several interesting avenues for future research come to the forefront. Firstly, an exciting direction would involve exploring the potential of our novel quantum machine learning models when combined with state-of-the-art policy gradient methods or actor-critic methods like DDPG [123], PPO [172], or A3C [138]. Such investigations can unveil synergies between classical reinforcement learning techniques and quantum enhancements, potentially leading to superior performance in complex learning tasks. Furthermore, it would be worthwhile to delve deeper into the capabilities of our novel quantum machine learning models as we scale up the number of available qubits.

The results of Chapter 6, which delve into the exponential separations between quantum and classical learning separations, prompt us to consider several interesting directions for future inquiry. Firstly, an interesting direction of future research would involve extending our investigations beyond the PAC learning framework. Delving into alternative learning frameworks, such as the Angluin learning framework [19], can broaden our understanding of quantum versus classical learning separations. Moreover, the outcomes of our work underscore the significance of probing the heuristic hardness of BQP-complete problems. Such investigations could give rise to novel learning separations within specific physics-inspired scenarios, thus contributing to a deeper understanding of the quantum advantage in practical applications.