



Universiteit
Leiden
The Netherlands

Capturing dynamics with noisy quantum computers

Dechant, D.S.

Citation

Dechant, D. S. (2026, February 17). *Capturing dynamics with noisy quantum computers*. Retrieved from <https://hdl.handle.net/1887/4290771>

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/1887/4290771>

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

With this chapter, we conclude the thesis. We restate and conclude the research questions outlined in Sec. 1.4 in the beginning of this thesis. Finally, we re-outline promising directions of future work.

6.1 Research overview

This thesis dealt with several aspects of capturing dynamics with noisy quantum computers. We introduced several concepts as well as the research questions in Chapter 1, which we explored in Chapters 2, 3, 4, and 5. While we drew conclusions at the end of each of Chapters 2, 3, 4, and 5, we now come back to the specific research questions stated in the beginning of the thesis and reflect on them based on the results derived in the earlier chapters.

The first two research questions are concerned with theoretical properties of quantum machine learning:

Research Question 1: *Can parameterized quantum circuits approximate functions as well as their derivatives?*

In Chapter 2, which is based on the previously published work in Ref. [88], we answered this question. We showed that parameterized quantum circuits can not only approximate square integrable functions arbitrarily close in the L^2 distance, but also other function classes and with respect to other distances. In particular,

by answering this research question, we proved that parameterized quantum circuits can approximate functions in the Sobolev space H^k with arbitrary precision under the Sobolev distance. Those spaces contain square-integrable functions whose partial derivatives up to order k are also square-integrable. The corresponding Sobolev distance is defined as the sum of the L^2 distances of the function and its partial derivatives up to order k . However, parameterized quantum circuits can only do so if certain conditions on the data input are met, for which we also provide a trivial solution.

Research Question 2: *How does an augmentation of the training data with derivatives of the target function influence generalization in quantum machine learning with parameterized quantum circuits?*

We also answered this question in Chapter 2, which is based on the results published in Ref. [88]. We proved a generalization bound that shows that a generalization of the approximation of both the function and its derivatives is possible if the training data includes not only labels of the target function, but also of its derivatives. Furthermore, we proved that including data of the derivatives of the target function also guarantees generalization bounds for the supremum and L^p distances. We found that the higher the dimension of the function, the more higher-order derivatives are required in order to achieve these bounds. These results give a theoretical explanation of earlier numerical findings that suggested improved generalization with classical neural networks [89], and thus also impact classical machine learning.

In the following research question, we focused on variational quantum algorithms for differential equations solving:

Research Question 3: *What is the total error arising in variational quantum algorithms for solving differential equations based on Runge-Kutta methods and which Runge-Kutta order minimizes the number of circuit evaluations needed?*

In Chapter 3, we provided an extensive error analysis and determined the resource requirements needed to achieve specific target errors, based on results published in Ref. [46]. In particular, we derived analytical error and resource estimates for scenarios with and without shot noise, examining shot noise in quantum measurements and truncation errors in Runge-Kutta methods. Our analysis did not take into account representation errors and hardware noise, as these are specific to the instance and the used device. We evaluated the implications of our results by applying them to two scenarios: classically solving a 1D ordinary differential equation and solving an option pricing linear partial differential equation with the variational algorithm, showing that the most resource-efficient methods are of order 4 and 2, respectively. We showed that

even those most resource-efficient methods require a number of shots multiple orders of magnitude higher than what seems to be feasible on near-term quantum computers. However, those resource estimates might be lower in practice, as several error bounds on which these estimates are based on are not tight. Although our study minimizes resources for the upper bound we derive, we hope that the resulting prescription is a good heuristic for allocating resources in practice. The results may also be of interest to the numerical analysis community as they involve the accumulation of errors in the function description, a topic that has hardly been explored even in the context of classical differential equation solvers.

In many variational quantum algorithms, estimating reduced density matrices (RDMs) of the quantum system forms an important subroutine. However, this task is challenging, among others due to the effects of shot noise stemming from a limited number of measurements. Our following research questions therefore asked:

Research Question 4: *Is it possible to mitigate the effects of shot noise in the quantum state tomography of reduced density matrices by enforcing physicality conditions organized as semidefinite programs?*

In Chapter 4, we answered this question, which is addressed in Ref. [90]. We proposed a method to mitigate shot noise by reinforcing certain physicality constraints on RDMs. The first kind of these constraints, which we called the enhanced-compatibility, require RDMs to be compatible with higher-dimensional RDMs. Secondly, we included constraints that we called overlapping-compatibility which enforce that overlapping RDMs are consistent on those subsystems on which they overlap. We organized these compatibility constraints in semidefinite programs to reconstruct RDMs from simulated data. Our approach yields, on average, tighter bounds for the same number of measurements compared to tomography without compatibility constraints. We further demonstrated the versatility and efficacy of our method by integrating it into an algorithmic cooling procedure to prepare low-energy states of local Hamiltonians.

In the last research question of this thesis, we explored quantum approaches for generative modeling of dynamical systems of the financial market. In particular, we stated:

Research Question 5: *Can quantum generative adversarial networks capture the distribution and stylized facts of financial time series on a qualitative level?*

In Chapter 5, we presented our answer to this question, based on the results shown in Ref. [91]. We trained quantum generative adversarial networks

(QGANs), composed of a quantum generator and a classical discriminator, and used two classical simulation approaches for the quantum generator: a full simulation of the quantum circuits, and an approximate simulation of the latter using matrix product states. We tested the effect of the choice of circuit depths and bond dimensions of the matrix product state simulation on the generated time series. Overall, the QGANs were generally successful in capturing most of the temporal correlations observed in real financial data. But depending on the hyperparameters of the model, the generated synthetic financial time series differed in how well they reproduced the properties of financial time series. These differences allow selecting a model that best reproduces the desired properties of financial time series for specific applications.

6.2 Future work

Throughout this thesis, we suggested several directions in which the results of the thesis can be extended, at the end of each of Chapters 2, 3, 4, and 5. Let us summarize here the most important ones.

The exploration of the expressivity and generalization of parameterized quantum circuits in Chapter 2 motivates further study of their impact on applications such as differential equation solving and option pricing. Further, it might be promising to examine the role of other distances for both the expressivity and generalization of parameterized quantum circuits.

In Chapter 3, we did not take into account representation errors and hardware noise, although we gave several ideas and explanations in how to incorporate them in practice. It is therefore natural to ask how to categorize these error sources in a more systematic manner, by using techniques such as in [131]. We determined the act of classically inverting an (in general ill-conditioned) matrix with noisy elements as the key difficulty in this set of algorithms. Exploring alternative ways for this step might therefore significantly improve the chances of applying these and related algorithms for real-world use cases. It might also be interesting to apply the ideas of Chapter 3 to classical differential solvers, as many of those are dealing with noisy data.

The method presented in Chapter 4 can be extended by including additional constraints, such as entropy constraints, into the SDP formulation to further refine the reconstruction process [195]. An interesting question is to explore its effect with other systems, such as frustration-free Hamiltonians or for quantum chemistry calculations, as in [207, 208]. Furthermore, it might be promising to combine our method with noise mitigation strategies and to apply it in the optimization of other local observables, such as correlation functions.

The quantum generative adversarial networks of Chapter 5 can possibly be applied as subroutines for applications such as option pricing [275] and risk analysis [276]. Moreover, a possible extension of our method is to train the model

to replicate correlated stocks of the S&P 500 index, motivated by research in community detection [277]. We gave several ideas for this extension in Chapter 5. We lastly suggest to explore possible improvement of the training of our model by several ways. In particular, as the model is trained with Wasserstein loss functions (see Eqs. (5.9) and (5.11)) that are taking the distribution of the time series into account, but not the temporal effects, an adaptation of the training to consider them as well might lead to a better recovery of those temporal effects.