



Universiteit
Leiden
The Netherlands

Quantum methods for machine learning and classical dynamics

Barthe, A.M.

Citation

Barthe, A. M. (2026, March 20). *Quantum methods for machine learning and classical dynamics*. Retrieved from <https://hdl.handle.net/1887/4297482>

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/1887/4297482>

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

Quantum computing offers a new model of computation. The goal is to use the principles of quantum mechanics to perform tasks that are expected to be intractable for classical computers. This thesis investigates how its development could impact learning algorithms and the simulation of classical physical systems. The overarching goal is to identify settings in which quantum computers provide a *provable* computational or learning advantage.

The first chapter introduces this work within the context of quantum computing and introduces three guiding questions. The first concerns the minimal resources required for universality in parameterized quantum models; the second asks how to construct provable quantum learning advantages; and the third examines the computational complexity of bosonic systems and their relation to classical and qubit-based computation.

The second chapter constitutes background on the notions of quantum computation, statistical learning theory, bosonic Hamiltonians, parameterized quantum circuits, and quantum complexity classes.

In Chapter 3, we study *quantum re-uploading models* (QRUs), i.e., parameterized circuits that repeatedly encode the input data. We analyze their expressivity, looking at the frequency spectrum of the functions realized by QRUs: as the number of re-uploads L grows, the average spectrum concentrates into a Gaussian with width scaling like $\sqrt{L}$, while the support grows like L . Practically, QRUs present an inductive bias for smooth functions, which is good for generalization but are limited in capturing fine details.

In Chapter 4, we move from supervised learning to generative modeling and introduce *expectation value samplers* (EVS): given a circuit $U(x)$ with random classical inputs x and a set of observables, the output is the vector of their expectation values. We prove the universality of the EVS model

on the hypercube $[-1, 1]^M$ in two regimes: (i) with observables of constant spectral norm on M qubits (ii) with $\Theta(\log M)$ qubits and $\Theta(M)$ -norm observables. These theorems make explicit the resource trade-off between qubit count and observable norm, which directly translates into shot complexity. We also show how EVS connects naturally to a many-body physics setting, highlighting when the model is a good fit for physical tasks.

In Chapter 5, we give a supervised learning task where a quantum learner provably outperforms any classical learner under standard complexity assumptions. The key ingredient is a *Fourier coefficient extraction* routine that reads spectral components of PQC-induced functions. This yields a feature map that is hard to compute classically, leading to a PAC learning separation. By Trotterization, we translate the same idea to *learning Hamiltonian dynamics* from input–output pairs: efficient on a quantum device, hard classically. This feature map also yields a kernel whose Gram matrix can be evaluated efficiently on a quantum computer.

In Chapter 6, we design a continuous-variable (CV) quantum algorithm for ordinary differential equations within the Koopman–von Neumann (KvN) framework. KvN embeds a classical ODE as a Hamiltonian evolution on an infinite-dimensional Hilbert space. This framework is naturally suited for *initial distribution problems*, that is, the evolution of a whole probability density rather than single-trajectory initial value problems. The chapter derives the structure constants of a relevant Lie algebra. Using higher-order Trotter formulas, the Hamiltonian evolution is compiled using just three gate types. We show how uncertainty in initial conditions translates into the squeezing level of the initial bosonic state.

In Chapter 7, we study the complexity of simulating large Gaussian bosonic circuits. Restricting to a simple family of interferometers yields a decision problem that we prove is BQP-complete. Adding structured squeezing boosts the computational power to PostBQP-hardness. Intuitively, squeezing plus simple interferometry can emulate postselection-like effects on expectation values, which explains the jump in complexity. This chapter further strengthens our understanding of how the computational complexity of simulating exponentially large classical systems relates to that of quantum computation.