

Provable learning separation for predicting time-evolution of quantum many-body systems

Rahul Bandyopadhyay,¹ Riccardo Molteni,¹ Jens Eisert,^{2,3,4} Vedran Dunjko,¹ and Sofiene Jerbi^{2,4}

¹ *$\langle aQa^L \rangle$, Applied Quantum Algorithms, Leiden University, Leiden, The Netherlands*

²*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany*

³*Fraunhofer Heinrich Hertz Institute, Berlin, Germany*

⁴*Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany*

Our work may be found under Ref. [1].

I. INTRODUCTION

It is widely believed that quantum computers are capable of simulating certain many-body systems more efficiently than classical computers [2, 3]. Guided by the intuition that similar advantages may hold for *learning* tasks as well, an immediate question arises: can one identify quantum machine learning (QML) separations for *physically motivated* learning tasks? An example of such a setting is the problem of predicting ground state properties of quantum systems, a quintessentially quantum learning task for which one may expect quantum learners to yield an advantage over classical ML. Within the natural regime of gapped phases of matter, however, the aforementioned learning problem has been found to be classically tractable [4–7]. Therefore, in this work, we instead consider the complementary setting of predicting *Hamiltonian time-dynamics* and show that this task exhibits a provable learning separation, premised on the well-founded complexity theoretic conjecture $\text{BQP} \not\subseteq \text{P/poly}$.

Related work. Our work is in a similar vein as Refs. [8, 9], which laid out a systematic framework for constructing learning separations via the BQP-completeness of the data-generating function, but differs from those works in that it applies that framework to learning unknown Hamiltonian dynamics, as opposed to an unknown *observable*, as in Ref. [9]. Prior to us, Ref. [10] also considered the task of predicting time-dynamics generated by an unknown Hamiltonian from classical measurement data, for which the authors of that work showed a learning separation as well. However, their learning problem is formulated in terms of a relatively restrictive family of Hamiltonians, namely, ones parametrized by $\mathcal{O}(\log n)$ -dimensional coefficient vectors. Consequently, this leads to an ML task that, while being physically motivated, does not capture the learning of local Hamiltonian dynamics in broad generality. Moreover, premised on the results of Ref. [11] and for the particular form of their training data, they provide complexity-theoretic arguments against the possibility of significant improvements of their method. In this work, by assuming access to a stronger form of training data, we avoid this obstruction and address the problem of predicting the dynamics of a more general family of Hamiltonians, namely, ones that encompass geometrically local systems.

Significance of present work. A pressing issue in QML remains the identification of realistic ML problems in many-body and condensed matter physics that exhibit quantum learning advantages. Our work makes tangible progress on this front, by demonstrating quantum advantage for a physical ML task, distinctly in more general regimes than those considered in prior work. Methodologically, we build connections between learning theory, quantum simulation, QML, and Hamiltonian complexity, and show how tools from these fields can be used in conjunction to explore quantum-classical learning separations, which we believe will be of strong interest to the QTML community. Our work also invites further investigation into learning problems of interest in various areas of physics, as well as the potential of quantum computers in solving these tasks more efficiently than classical computers.

A. Contributions

Concretely, our contributions are as follows:

1. We formulate a supervised learning task where the goal is to predict expectation values of quantum states that have been time-evolved by some unknown, *low-intersection* n -qubit Hamiltonian of the form $H(\lambda) := \sum_{a \in [M]} \lambda_a E_a$, where $\lambda := (\lambda_a)_{a \in [M]} \in [-1, 1]^M$ and $(E_a)_{a \in [M]}$ are unique, traceless, non-identity Pauli strings. We call $H(\lambda)$ low-intersection if each qubit is acted upon non-trivially by a number of Pauli strings that is at most $\mathcal{O}(1)$, leading to a number of terms scaling as $M \in \mathcal{O}(n)$, thereby encompassing many-body systems of the kind encountered in condensed matter physics.
2. We establish a learning separation for the aforementioned learning problem, premised on the assumption $\text{BQP} \not\subseteq \text{P/poly}$. In particular, we provide an efficient quantum learning algorithm whose training phase is based on a simplified variant of the Hamiltonian learning protocol of Haah, Kothari, and Tang [12] and the deployment stage of which makes use of Hamiltonian simulation. On the other hand, to prove classical hardness, we rely on the construction of Oliveira




Figure 1. (a) Training samples are generated by the time-evolution of a low-intersection Hamiltonian $H(\lambda)$ for randomly chosen times $t \sim \text{Unif}([0, T])$ acting on some initial probe state $|\psi(x)\rangle$. (b) We consider as our labels a vector of expectation values evaluated on a set of certain pre-determined observables. This has been schematically depicted in the above figure as a collection of measurement samples. (c) We filter out the short-time samples from our training data and use them to identify the underlying data-generating Hamiltonian via a simple Hamiltonian learning algorithm that is classical. (d) Finally, we use a quantum computer to successfully perform inference on a newly given data point. By universality of Hamiltonian evolution, this can involve predicting a BQP-complete computation, which is believed to be hard for classical learners.

and Terhal [13], which gives a low-intersection version of the Feynman-Kitaev clock Hamiltonian whose time-evolution realizes a BQP-complete computation for polynomially large evolution times.

II. LEARNING PROBLEM

At a high level, the goal of our learning task is to predict, for various given times and input states, a classical description of these states after they have been time evolved under an unknown Hamiltonian. The training data for this task is generated as follows: for different inputs $x \in \{0, 1\}^{\text{poly}(n)}$ and evolution times $t \in [0, T]$, an n -qubit input state vector $|\psi(x)\rangle$ is prepared and time-evolved by a low-intersection Hamiltonian $H(\lambda)$ for time t , i.e., it evolves by $U(t, \lambda) = e^{-iH(\lambda)t}$. The classical description of this time-evolved state then takes the form of expectation values of a certain pre-determined set of observables that are gathered via repeated preparations and measurements. The Hamiltonian parameters $\lambda \in [-1, 1]^M$ are fixed but unknown, with $M \in \mathcal{O}(n)$, which follows from the low-intersection property.

Definition 1 (Concept class). Let $\mathcal{H} := \{H(\lambda) \mid \lambda \in [-1, 1]^M\}$ be a family of n -qubit low-intersection Hamiltonians. Let $x \in \{0, 1\}^{\text{poly}(n)}$ and $\{Q_a\}_{a \in [M]}$ be a set of certain observables. Then, we define the concept class

$$\mathcal{C}_{U_{\text{enc}}, T}^{\mathcal{H}} := \left\{ c_\lambda : \{0, 1\}^{\text{poly}(n)} \times [0, T] \ni (x, t) \mapsto \left(\langle \psi(x) | U^\dagger(t, \lambda) Q_a U(t, \lambda) | \psi(x) \rangle \right)_{a \in [M]} \in \mathbb{R}^M \mid \lambda \in [-1, 1]^M \right\}, \quad (1)$$

where $U(t, \lambda) := e^{-iH(\lambda)t}$, such that $|\psi(x)\rangle = U_{\text{enc}}(x) |0\rangle^{\otimes n}$, where U_{enc} is a data-encoding unitary and $T \in \mathcal{O}(\text{poly}(n))$. For some input pair $(x^{(j)}, t^{(j)})$, we define the action of the encoding unitary U_{enc} as preparing a product stabilizer probe state when $x_1^{(j)} = 0$, or a computational basis state when $x_2^{(j)} = 1$.

Given training data of the form $\mathcal{T}^\lambda := \{(x^{(j)}, t^{(j)}, y^{(j)})\}_{j=1}^N$, where $(x^{(j)}, t^{(j)}) \sim \mathcal{D}$ for $\mathcal{D}$ belonging to a family of certain allowed distributions and $y^{(j)} := c_\lambda(x^{(j)}, t^{(j)})$, the goal of the ML algorithm is to learn a model $h(x, t)$ that approximates the unknown concept generating the data. This is formally stated as the following.

Definition 2 (Efficient PAC-learning condition). *The concept class $\mathcal{C}_{U_{\text{enc}}, T}^H$ is PAC-learnable with respect to a set of distributions $\{\mathcal{D}_i\}_i$ over $(x, t) \in \{0, 1\}^{p(n)} \times [0, T]$ if there exists an algorithm $\mathcal{A}$ such that, for all $\varepsilon, \delta > 0$, for any c_λ in $\mathcal{C}_{U_{\text{enc}}, T}^H$, and any input distribution $\mathcal{D}_i \in \{\mathcal{D}_i\}_i$, $\mathcal{A}$ receives as input a training set $\mathcal{T}^\lambda$ of size $\mathcal{O}(\text{poly}(n, 1/\varepsilon, 1/\delta))$ and outputs a hypothesis $h(\cdot) = \mathcal{A}(\mathcal{T}^\lambda, \varepsilon, \delta, \cdot)$ which satisfies*

$$\mathbb{E}_{(x,t) \sim \mathcal{D}_i} [\|c_\lambda(x, t) - h(x, t)\|_\infty] \leq \varepsilon \quad (2)$$

with probability at least $1 - \delta$. The concept class is said to be efficiently PAC-learnable if the time-complexity of $\mathcal{A}$ is $\mathcal{O}(\text{poly}(n, \varepsilon^{-1}, \delta^{-1}))$.

III. QUANTUM LEARNABILITY

Theorem 1 (Quantum learnability, informal). *There exists a family of distributions such that, with respect to each distribution in that family, the concept class in Def. 1 is efficiently PAC-learnable using a quantum learning algorithm $\mathcal{A}_Q$. That is, the sample- and time-complexity of $\mathcal{A}_Q$ is $\mathcal{O}(\text{poly}(n, 1/\varepsilon, 1/\delta))$ for precision and confidence parameters $\varepsilon, \delta > 0$.*

Proof-sketch. Our analysis of quantum learnability proceeds by considering the expansion of the following expectation value:

$$\mathcal{F}_{a,t} := \text{Tr} \left[Q_a e^{-iH(\lambda)t} P_a e^{iH(\lambda)t} \right] = 4\lambda_a t + \mathcal{O}(t^2), \quad (3)$$

where, per the convention of Ref. [12], P_a is any single-qubit Pauli that anti-commutes with the Hamiltonian term E_a , and $Q_a := i[P_a E_a] = 2iP_a E_a$. We then consider the map

$$\tilde{\mathcal{F}}_{t^*}: [-1, 1]^M \ni (\lambda_a)_{a \in [M]} \mapsto \left(\tilde{\mathcal{F}}_{b,t^*} := \frac{1}{t^*} \int_0^{t^*} \mathcal{F}_{b,t} dt \right)_{b \in [M]} \in [-1, 1]^M, \quad (4)$$

$$\text{where } t^* \text{ is such that } \frac{1}{t^*} \int_0^{t^*} \mathcal{F}_{b,t} dt = 2\lambda_a t^* + \underbrace{\dots}_{\leq \varepsilon_{\text{tail}}} \quad (5)$$

for $\varepsilon_{\text{tail}}$ some allocated error budget. Algorithmically, we will compute the time-averaged integral via Monte-Carlo estimation, which we perform simply by taking the empirical average over training samples contained in the short-time window $[0, t^*] \subset [0, T]$. Dividing the empirical average by $2t^*$ gives us an estimate $\hat{\lambda}_a$ of λ_a , up to some precision controlled by $\varepsilon_{\text{tail}}$ and the accuracy of Monte-Carlo estimation. We then calculate the sample complexity of short-time samples, which we denote by N_{t^*} . We do this by bounding the difference between the Monte-Carlo estimate and the time-averaged integral, which we show to be the expectation of the random variable from which the training samples are drawn, conditioned on $t \in [0, t^*]$, thereby allowing us to use Hoeffding's inequality. Finally, we guarantee that enough training samples are contained in the short-time window, so as to ensure efficient execution of the training algorithm by using a Chernoff bound to compute the total sample complexity N , incurring only a polynomial blow-up with respect to N_{t^*} . The sample complexity of the inference stage is determined by the cost of implementing the classical shadows protocol, which is logarithmic in the number of observables being predicted and the inverse of the failure probability, and inverse-polynomial in the desired precision. We also find that the overall runtime of our learning procedure is polynomial in the relevant parameters. $\square$

IV. CLASSICAL HARDNESS

Theorem 2 (Classical hardness, informal). *For any (Promise)BQP-complete language, there exists a Hamiltonian family H_{BQP} specifying a concept class $\mathcal{C}_{U_{\text{enc}}, T}^{H_{\text{BQP}}}$ as per Def. 1 and a family of input distributions such that no randomized polynomial-time classical algorithm $\mathcal{A}_C$ satisfies the PAC-learning condition of Def. 2 for this concept class, unless $\text{BQP} \subseteq \text{P/poly}$.*

Proof-sketch. Our proof of classical hardness relies primarily on the following two ingredients: first, a Hamiltonian family $\{H_{\text{BQP}}(\lambda)\}_\lambda$ due to Oliveira and Terhal [13], which gives a version of the Feynman-Kitaev clock Hamiltonian that satisfies the low-intersection property, so as to have quantum learnability for the classically hard instance. We rigorously show how the long-time dynamics of this Hamiltonian family can simulate BQP-complete computations with success probability arbitrarily close to 1. Secondly, towards a proof by contradiction, we make use of the seminal boosting result of Schapire that allows us to promote any efficient classical PAC-learner under our choice of input distributions to polynomial-size classical circuits for deciding BQP languages, which in turn would imply $\text{BQP} \subseteq \text{P/poly}$ [14]. $\square$

-
- [1] R. Bandyopadhyay, R. Molteni, J. Eisert, V. Dunjko, and S. Jerbi, [Provable learning separation for predicting time-evolution of quantum many-body systems](#) (2026), [arXiv:2607.06472 \[quant-ph\]](#).
 - [2] R. P. Feynman, in *Feynman and computation* (cRc Press, 2018) pp. 133–153.
 - [3] S. Lloyd, *Science* **273**, 1073 (1996).
 - [4] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, *Science* **377**, [10.1126/science.abk3333](#) (2022).
 - [5] L. Lewis, H.-Y. Huang, V. T. Tran, S. Lehner, R. Kueng, and J. Preskill, *Nature Comm.* **15**, 895 (2024).
 - [6] M. Wanner, L. Lewis, C. Bhattacharyya, D. Dubhashi, and A. Gheorghiu, [Predicting ground state properties: Constant sample complexity and deep learning algorithms](#) (2024), [arXiv:2405.18489 \[quant-ph\]](#).
 - [7] C. Rouzé, D. Stilck França, E. Onorati, and J. D. Watson, *Nature Comm.* **15**, 7755 (2024).
 - [8] C. Gyurik and V. Dunjko, [arXiv preprint arXiv:2306.16028](#) (2023).
 - [9] R. Molteni, C. Gyurik, and V. Dunjko, [Exponential quantum advantages in learning quantum observables from classical data](#) (2024), [arXiv:2405.02027 \[quant-ph\]](#).
 - [10] A. Barthe, M. Y. Rad, M. Grossi, and V. Dunjko, [Quantum advantage in learning quantum dynamics via fourier coefficient extraction](#) (2025), [arXiv:2506.17089](#).
 - [11] S. Arunachalam, A. B. Grilo, and A. Sundaram, [Quantum hardness of learning shallow classical circuits](#) (2019), [arXiv:1903.02840 \[quant-ph\]](#).
 - [12] J. Haah, R. Kothari, and E. Tang, *Nature Physics* **20**, 1027–1031 (2024).
 - [13] R. Oliveira and B. M. Terhal, [The complexity of quantum spin systems on a two-dimensional square lattice](#) (2008), [arXiv:quant-ph/0504050 \[quant-ph\]](#).
 - [14] R. E. Schapire, *Machine learning* **5**, 197 (1990).