

Quantum Science and Technology



PAPER

OPEN ACCESS

RECEIVED
3 October 2021REVISED
1 January 2022ACCEPTED FOR PUBLICATION
26 January 2022PUBLISHED
1 March 2022

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Sample complexity of learning parametric quantum circuits

Haoyuan Cai¹ , Qi Ye¹ and Dong-Ling Deng^{1,2,*} ¹ Center for Quantum Information, IIIS, Tsinghua University, Beijing 100084, People's Republic of China² Shanghai Qi Zhi Institute, 41th Floor, AI Tower, No. 701 Yunjin Road, Xuhui District, Shanghai 200232, People's Republic of China

* Author to whom any correspondence should be addressed.

E-mail: dldeng@tsinghua.edu.cn**Keywords:** quantum machine learning, sample complexity, PAC learnable, quantum circuit

Abstract

Quantum computers hold unprecedented potentials for machine learning applications. Here, we prove that physical quantum circuits are probably approximately correct learnable on a quantum computer via empirical risk minimization: to learn a parametric quantum circuit with at most n^c gates and each gate acting on a constant number of qubits, the sample complexity is bounded by $\tilde{O}(n^{c+1})$. In particular, we explicitly construct a family of variational quantum circuits with $O(n^{c+1})$ elementary gates arranged in a fixed pattern, which can represent all physical quantum circuits consisting of at most n^c elementary gates. Our results provide a valuable guide for quantum machine learning in both theory and practice.

1. Introduction

Over the past few decades, machine learning, especially deep learning, has made dramatic progress [1, 2] in a wide range of tasks, such as playing the game of Go [3, 4], protein structure prediction [5], and computer vision [6], etc. More recently, the interplay between machine learning and quantum physics has attracted tremendous interest [7–11], giving birth to an emergent research Frontier of quantum machine learning. A number of notable quantum algorithms, such as the Harrow–Hassidim–Lloyd (HHL) algorithm [12], quantum generative models [13], and quantum support vector machine [14], have been designed to enhance, speed up, or innovate machine learning with quantum devices. These algorithms bear the intriguing potentials of exhibiting exponential advantages compared to their classical counterparts, although subtle caveats do exist and require careful examinations in practice [15].

In 1984, Valiant introduced the probably approximately correct (PAC) learning model [16], which gives a complexity-theoretical foundation and a mathematically rigorous framework for studying machine learning. Since then, the PAC learning model has been extensively studied in various machine learning scenarios to understand why and when efficient learning is possible or not [17, 18]. With the rapid progress in quantum computing [19–21], practical applications of quantum machine learning have become more and more realistic [22–26]. A natural problem is then to generalize the PAC learning model to quantum learning scenarios. Indeed, notable progress has been made along this direction [27–37]. For example, in reference [28] Chung and Lin have studied the sample complexity of learning quantum channels and demonstrated that we can PAC-learn a polynomial-size quantum circuit with a polynomial number of samples. In addition, in reference [33] Bu *et al* investigated the Rademacher complexity of quantum circuits in the framework of quantum resource theories [38]. They introduced a resource measure of magic for quantum channels based on the (p, q) group norm and found useful bounds for how the statistical complexity scales with resources in the quantum circuits. Yet, this fledgling research direction is still in its rapidly growing early phase and many important issues remain barely explored.

In this paper, we study the problem of the sample complexity for learning parametric quantum circuits. We focus on the supervised learning scenarios and prove that all the unitary physical quantum circuits are PAC learnable on a quantum computer via empirical risk minimization. More concretely, we prove the following two theorems: (1) any physical n -qubit quantum circuit consisting of at most n^c unitary gates with each gate acting on a constant number of qubits can be represented in an exact fashion by a family of

variational quantum circuits with $O(n^{c+1})$ elementary gates arranged in a fixed uniform pattern; (2) this family of variational quantum circuits is PAC learnable. Since most quantum circuits that can be efficiently implemented on a quantum computer, such as the circuits for the Shor's algorithm [39] or the HHL algorithm [12], contain at most a polynomial number of gates, our results imply that they are all PAC learnable with a quantum computer.

2. Results

2.1. Notations and the general setting

We define the concept class \mathcal{C} as the collection of all the n -qubit parametric quantum circuits with at most n^c unitary gates, each gate acting on at most b qubits (b, c are constant numbers independent of n). We note that \mathcal{C} is general enough to include most quantum circuits in practical applications. Here, we study the learnability of the quantum circuits in \mathcal{C} under the PAC learning framework [18]. Let $C \in \mathcal{C}$ be any n -qubit circuit in this concept class. When we input an n -qubit pure state $|\psi\rangle_{\text{in}}$ to C , we will get an output n -qubit pure state $|\psi\rangle_{\text{out}} = C|\psi\rangle_{\text{in}}$. Therefore, C can be viewed as a function $f_C : \mathcal{X} \rightarrow \mathcal{Y}$, where its domain \mathcal{X} and range \mathcal{Y} are both the set of all n -qubit pure states. In this work, we write $x \in \mathcal{X}$ as an abbreviation of the n -qubit quantum state $|\psi(x)\rangle$, and similarly for $y \in \mathcal{Y}$. With these notations, we sometimes write $y = f_C(x)$ to denote $|\psi(y)\rangle = C|\psi(x)\rangle$ for simplicity.

We consider the supervised learning scenario [18] and denote the training set of size m as $S = \{(x_1, y_1), (x_2, y_2), \dots, (x_m, y_m)\}$. Under the PAC learning framework, to learn the unknown circuit C , assume we have m independent n -qubit input samples $\{x_1, x_2, \dots, x_m\} \in \mathcal{X}^m$, we can input them into C , and obtain output states $\{y_1, y_2, \dots, y_m\} \in \mathcal{Y}^m$. The essential task of supervised learning is then to learn from S a hypothesis function (here a quantum circuit F) that can approximate the target function $f_C(x)$. This might be accomplished by minimizing certain loss functions over a set of variational model parameters. More concretely, we construct a variational quantum circuit F consisting of multiple gates with some of them having tunable parameters. By tuning these parameters, we can use F to represent different functions $f_h : \mathcal{X} \rightarrow \mathcal{Y}$, and we define our hypothesis space \mathcal{F} as the collection of all the functions f_h that F can represent. Given m independent samples and a tunable quantum circuit F , we can use the following process to make our circuit F become a good approximation of C . We can tune the parameters of F according to the training set S , so that when we put the state x_i ($i = 1, 2, \dots, m$) into the input of F , the output of F will be a good approximation of y_i . By the PAC learning theory, assuming that \mathcal{F} has good generalization power, f_h 's decent performance on the training set can imply its good performance over the whole sample space.

The effectiveness of the above process is based on two assumptions. First, the space of \mathcal{F} should be large enough, so that given any set of samples $S = (x_i, y_i)_{i=1,2,\dots,m}$, we can always find a function $f_h \in \mathcal{F}$, such that $f_h(x_i)$ can approximate y_i with a small error for all $i = 1, 2, \dots, m$. Second, the space of \mathcal{F} should not be too large or complex, so that \mathcal{F} has favorable generalization power to generalize its performance from the training set S to the true probability distribution that S is sampled from. This is a reflection of the Occam's razor principle [18]. Therefore, we need to design a variational quantum circuit class \mathcal{F} , which meets the following two requirements simultaneously in order to learn f_C :

- **R1:** For any $C \in \mathcal{C}$, there exists a hypothesis function $f_h \in \mathcal{F}$, such that $f_h(x) = f_C(x)$ for all $x \in \mathcal{X}$.
- **R2:** The hypothesis space \mathcal{F} satisfies the PAC learnability.

We note that the first requirement **R1** is stronger than the first assumption, because the function f_h in **R1** is the same as f_C . Hence, the training error f_h on the set of samples S is necessarily zero, whereas the first assumption only requires that f_h has a small training error for S . In supervised learning, obtaining high-quality training samples is usually resource-demanding in practice. Thus, studying the sample complexity becomes crucial. In the following, we will study the sample complexity of learning parametric quantum circuits and rigorously prove that any physical quantum circuit is PAC learnable.

2.2. A family of universal variational circuits

To meet **R1**, we should ensure that F has representation power for all the n -qubit parametric quantum circuits C in \mathcal{C} . We observe that any quantum circuit C can be decomposed as a sequence of $O(n^c)$ number of H gates, $R_x(\theta)$ gates, and CNOT gates (see the proposition 1 in appendix A). Thus in our construction of F , we also use these three kinds of gates and arrange them into a uniform pattern, so that its scaling to quantum circuits with more qubits is clear. The construction of F is illustrated in figure 1, which is based on block assembling, i.e., assembling some relatively small gadgets to form a more complicated block. By convention, we denote the three Pauli matrices by X , Y , and Z . A well-known result about quantum circuit states that any single-qubit unitary gate can be expressed as $e^{i\alpha} R_z(\beta) R_x(\gamma) R_z(\delta)$, where α, β, γ , and δ are

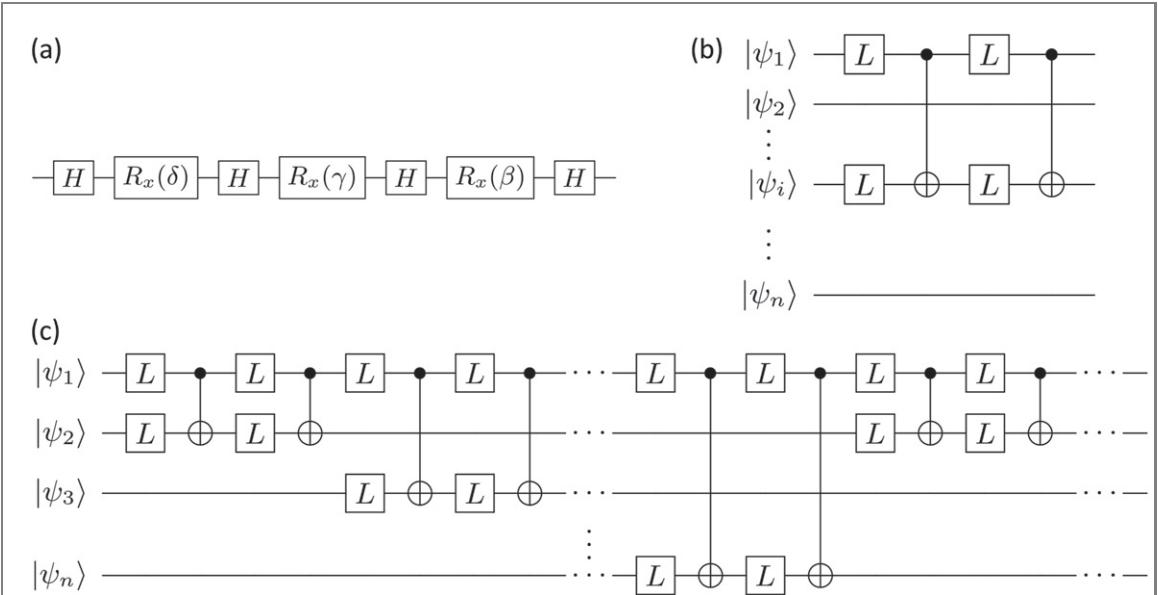


Figure 1. Pictorial illustration for the construction of the hypothesis quantum circuits. (a) The elementary (level-1) block L . This block contains four Hadamard gates and three single-qubit rotations along the x direction with rotation angles parameterized by δ , γ , and β , respectively. (b) The level-2 block B_i constructed based on four level-1 block and two CNOT gates with the first qubit being the control qubit and the i th qubit being the target one. (c) The constructed hypothesis quantum circuit, which consists of Mn^c repeated layers of $B_n B_{n-1} \dots B_2$.

four real numbers, and $R_z(\theta) = e^{-i\theta Z/2}$ and $R_x(\theta) = e^{-i\theta X/2}$ are the rotation operators along the z -axis and x -axis on the Bloch sphere, respectively [40]. Inspired by this, we define a set of basic gadgets (we call them level-1 blocks) to be $L = \{HR_x(\beta)HR_x(\gamma)HR_x(\delta)H|\beta, \gamma, \delta \in [0, 2\pi]\}$. In this way, we can tune the parameters of a level-1 block so that it can represent all the single-qubit unitary gates up to an irrelevant global phase factor $e^{i\alpha}$. In addition, when we set $\beta = \gamma = \delta = 0$, the level-1 block will reduce to the identity gate.

Using level-1 blocks we can construct level-2 blocks B_i ($i = 2, 3, \dots, n$). First we put two level-1 blocks (denoted as L) at qubit 1 and qubit i , and then insert two CNOT_{1*i*} gates, as shown in figure 1(b). Here, we use CNOT_{1*i*} to denote the controlled-NOT gate between the first and i th qubits, with the first qubit being the control qubit and the i th one being the target one. With this, the desired hypothesis quantum circuit F can be constructed as:

$$F = (B_n B_{n-1} \dots B_2)^{Mn^c}, \quad (1)$$

where M is a constant independent of the number of qubits n . We mention that the hypothesis circuit F has a uniform structure for arbitrary system sizes. In addition, only the x -rotations contain variational parameters, and it is straightforward to obtain that the total number of parameters used for defining F scales as $O(n^{c+1})$.

For an arbitrary quantum circuit U , we say that it can be represented by the variational circuit F if there exists a solution to the parameters (denoted collectively as θ) in F such that $F = U$ up to an irrelevant global phase. Now, we are ready to give the first theorem stating that for any $C \in \mathcal{C}$, we can use F to represent it:

Theorem 1. *For any $C \in \mathcal{C}$, there exists a hypothesis function $f_h \in \mathcal{F}$, such that $f_h(x) = f_C(x)$ for all $x \in \mathcal{X}$.*

Proof. We first give a high-level intuition for the proof. We note that any gate acting on a constant number of qubits can be decomposed into a quantum circuit with a constant number of elementary gates, namely the CNOT gates, H gates, and $R_x(\theta)$ gates. Thus, any $C \in \mathcal{C}$ can be decomposed into a quantum circuit with $O(n^c)$ elementary gates. In addition, as our hypothesis quantum circuit F consists of Mn^c layers and every two layers can represent one arbitrary elementary gate acting on any pair of qubits through tuning the parameters properly, we can prove that C can be represented by F in an exact fashion.

The complete proof is as follows. First, we note that C consists of $O(n^c)$ elementary gates by the definition of \mathcal{C} . By proposition 1 in appendix A, the circuit C can be written as a product form $U_l U_{l-1} \dots U_1$, where $l = O(n^c)$, and each U_i is either CNOT_{1*i*} or a single-qubit unitary gate V_j on qubit j .

Denoting $B_n B_{n-1} \dots B_2$ as one layer of level-2 blocks, then we define a block string of d layers as follows:

$$(B_n B_{n-1} \dots B_2)^d.$$

We define T as the minimal number so that a block string of T layers can represent $U_l U_{l-1} \dots U_1$. As our hypothesis circuit F contains Mn^c layers in total, we will show that to represent $U_l U_{l-1} \dots U_1$, the minimal number of layers needed in the block string is no greater than Mn^c . Then, when the previous T layers in F have represented C exactly, by the first part of proposition 2 in appendix A, we can set the remaining $(Mn^c - T)$ layers to be the identity gate on the n qubits. Therefore, we can show that F can represent C in an exact fashion and complete the proof of theorem 1.

To show that $T \leq Mn^c$, we will prove that for each $U_i, i = 1, 2, \dots, l$, we only need two layers to represent U_i exactly. Then putting them together, we can prove that we need only $2l$ layers to represent $U_l U_{l-1} \dots U_1$, which yields $T \leq 2l$.

We fix any $i \in \{1, 2, \dots, l\}$. By proposition 2, one layer $B_n B_{n-1} \dots B_2$ can represent any single-qubit unitary gate V_j acting on any qubit $j = 1, 2, \dots, n$ up to a global phase $e^{i\alpha}$. Moreover, two layers $(B_n B_{n-1} \dots B_2)^2$ can represent CNOT_{1j} for any $j = 2, 3, \dots, n$. We note that U_i is either a single-qubit gate V_j acting on some qubit j , or a two-qubit gate CNOT_{1j} . Therefore, we only need two layers at most to represent U_i exactly.

As we have shown that $T \leq 2l = O(n^c)$, by choosing a large enough constant M , we can prove that $T \leq Mn^c$ and complete the proof of the theorem. \square

Theorem 1 shows that given any quantum circuit $C \in \mathcal{C}$, there always exists a solution to the parameters, such that the circuit F can simulate the quantum circuit C acting on the n -qubits with zero error. Therefore, given any training set $S = (x_i, y_i)_{i=1,2,\dots,m}$ sampled independently from some distribution P over $\mathcal{X} \times \mathcal{Y}$, if we have $y_i = f_C(x_i)$ for all $i = 1, 2, \dots, m$, we can find an instance $f_h \in \mathcal{F}$ with zero training error. In fact, this theorem has a wider range of applications. When a quantum circuit consists of fewer than n^c gates, we can add some identity gates after it, so our theorem covers all the quantum circuits containing no more than n^c gates. In other words, we can use only $O(n^{c+1})$ gates, which are arranged in a uniform pattern, to represent all the circuits with n^c gates or fewer. We remark that the number of gates in many famous quantum circuits, such as quantum support vector machine [14], HHL algorithm [12], and quantum Fourier transform [41], scales polynomially with the number of qubits. Therefore, all of these circuits above can be represented exactly by our circuit F .

2.3. PAC learnability of \mathcal{F}

In the PAC setting, we usually assume that the input samples are randomly generated from certain unknown probability distribution. As a result, when the hypothesis space covers the underlying distribution and the training dataset is large enough, both the training and generalization error should be small. In this paper, our hypothesis space \mathcal{F} has been proved in theorem 1 to be able to cover all the parametric quantum circuits in \mathcal{C} . Now, we study the sample complexity for training a circuit $F \in \mathcal{F}$ to represent $C \in \mathcal{C}$. To this end, we define a measure of the distance between two pure states in \mathcal{Y} , which is used as the loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, 1]$. Specifically, we define the loss function $\mathcal{L}(y_1, y_2)$ to be the trace distance of two quantum states $|\psi(y_1)\rangle$ and $|\psi(y_2)\rangle$:

$$\mathcal{L}(y_1, y_2) = \frac{1}{2} \|\langle \psi(y_1) | \psi(y_1) \rangle - \langle \psi(y_2) | \psi(y_2) \rangle\|_1,$$

where $\|\cdot\|_1$ denotes the trace norm of a matrix. Given a hypothesis function $f_h \in \mathcal{F}$ and the training set $S = (x_i, y_i)_{i=1,2,\dots,m}$ sampled independently from some distribution P , we can define the empirical risk of f_h , which is also known as in-sample error:

$$\hat{R}(f_h) = \frac{1}{m} \sum_{i=1}^m \mathcal{L}(y_i, f_h(x_i)).$$

The risk of a hypothesis function $f_h \in \mathcal{F}$ is then defined as the average loss of f_h over the probability distribution P :

$$R(f_h) = \int_{\mathcal{X} \times \mathcal{Y}} \mathcal{L}(y, f_h(x)) dP(x, y).$$

Our goal is to find a hypothesis $f_h \in \mathcal{F}$ and minimize its risk $R(f_h)$. As the parametric quantum circuit C is a black box in our setting, we do not know the probability distribution P . In the learning process, we use the training set S and find an empirical risk minimizer $\hat{h} \in \mathcal{F}$ over S . For convenience, given the training set S and the probability distribution P , we define $\hat{h} = \arg \min_{h' \in \mathcal{F}} \hat{R}(h')$ to be the empirical risk minimizer, and

$h = \arg \min_{h' \in \mathcal{F}} R(h')$ to be the risk minimizer. Now we formally introduce the definition of the PAC learnability for completeness [42]:

Definition 1 (PAC learnability). A hypothesis space \mathcal{F} is PAC learnable, if there exists a function $\nu : (0, 1)^2 \rightarrow \mathbb{N}$, such that for all $(\epsilon, \delta) \in (0, 1)^2$ and all probability measures P over $\mathcal{X} \times \mathcal{Y}$, when the size of training set $|S| \geq \nu(\epsilon, \delta)$, we have

$$\mathbb{P}_S \left(R(\hat{h}) - R(h) \leq \epsilon \right) \geq 1 - \delta. \quad (2)$$

Here $\mathbb{P}_S(A)$ denotes the probability that event A happens over repeated sampling of the training set S .

We note that when we randomly select a state $x \in \mathcal{X}$, input it into the circuit C and get an output state $y = f_C(x) \in \mathcal{Y}$, the resulting probability distribution P of state pairs (x, y) will satisfy $R(h) = 0$, because we can find an instance $f_h \in \mathcal{F}$ equal to f_C by theorem 1. Therefore, if our \mathcal{F} is PAC learnable, after we prepare the training samples S and get an empirical minimizer \hat{h} , with probability $1 - \delta$, the average loss of \hat{h} over P will be no larger than ϵ , i.e., $R(\hat{h}) \leq \epsilon$. Now we are going to prove that \mathcal{F} is PAC learnable, and the sample complexity ν is polynomial with n , $1/\epsilon$, and $\ln \frac{1}{\delta}$.

Theorem 2. *The hypothesis space \mathcal{F} satisfies the PAC learnability, with sample complexity $\nu(\epsilon, \delta) = O\left(\frac{1}{\epsilon^2}(n^{c+1} \ln \frac{n}{\epsilon} + \ln \frac{1}{\delta})\right)$.*

Proof. The essential idea for the proof relies on the discretization of \mathcal{F} . First, we construct a finite set of hypothesis functions \mathcal{F}' , such that for each function $f_h \in \mathcal{F}$, we can find a function $f'_h \in \mathcal{F}'$ close enough to f_h . Then we use lemma 2 in appendix B to show that \mathcal{F}' is PAC learnable. Finally, for any $f_h \in \mathcal{F}$ and its corresponding f'_h , as f_h and f'_h are close enough, we can prove that their risk and empirical risk are close as well. Therefore, we obtain that \mathcal{F} is PAC learnable.

For clarity, we denote l as the total number of $R_x(\theta)$ gates in the circuit F , and observe that $l \leq 12Mn^{c+1}$. We recall that \mathcal{F} is defined as the collection of all the functions $f_h : \mathcal{X} \rightarrow \mathcal{Y}$ that the circuit F can represent by tuning the value of the parameters $\theta = (\theta_1, \theta_2, \dots, \theta_l)$, where $\theta_i \in [0, 2\pi)$ is the variational parameter characterizing the i th x -rotation. Now we define a finite set $\mathcal{F}' \subseteq \mathcal{F}$ in this way: \mathcal{F}' is the collection of all the functions $f'_h : \mathcal{X} \rightarrow \mathcal{Y}$ that circuit F can represent by tuning the value of all the θ_i in $\{0, e, 2e, \dots, Ne\}$, where $e = \frac{\epsilon}{6Kn^{c+1}}$, $N = \lfloor \frac{2\pi}{e} \rfloor$, and K is a large enough constant. As there are $l = O(n^{c+1})$ rotational gates in circuit F in total, we have

$$|\mathcal{F}'| = \left\lceil \frac{12\pi Kn^{c+1}}{\epsilon} \right\rceil^{O(n^{c+1})},$$

which is finite. As a result, we can plug \mathcal{F}' and $\epsilon' = \epsilon/6$ into lemma 2 to obtain that when $|S| \geq \frac{18}{\epsilon^2}(\ln |\mathcal{F}'| + \ln \frac{2}{\delta})$,

$$\mathbb{P}_S \left(\forall f'_h \in \mathcal{F}' : |\hat{R}(f'_h) - R(f'_h)| \leq \frac{\epsilon}{6} \right) \geq 1 - \delta. \quad (3)$$

We fix all the parameters θ in circuit F , and we will get an arbitrary hypothesis function $f_h \in \mathcal{F}$. Then we can round all the parameter θ of circuit F into their nearest multiples of e in $\{0, e, 2e, \dots, Ne\}$, and we will get a new hypothesis function $\tilde{f}_h \in \mathcal{F}'$. By proposition 5, we obtain that for any $f_h \in \mathcal{F}$,

$$|R(\tilde{f}_h) - R(f_h)| \leq \frac{\epsilon}{6}, \quad (4)$$

$$|\hat{R}(\tilde{f}_h) - \hat{R}(f_h)| \leq \frac{\epsilon}{6}. \quad (5)$$

Combining the three inequalities (3)–(5) together, we arrive at

$$\mathbb{P}_S \left(\forall f_h \in \mathcal{F} : |\hat{R}(f_h) - R(f_h)| \leq \frac{\epsilon}{2} \right) \geq 1 - \delta, \quad (6)$$

when $|S| \geq \frac{18}{\epsilon^2}(\ln |\mathcal{F}'| + \ln \frac{2}{\delta})$. To prove that \mathcal{F} is PAC learnable, we recall our notations that

$h = \arg \min_{h' \in \mathcal{F}} R(h')$, and $\hat{h} = \arg \min_{h' \in \mathcal{F}} \hat{R}(h')$. Combining the inequality (6) and proposition 6, we obtain that

$$\mathbb{P}_S \left(R(\hat{h}) - R(h) \leq \epsilon \right) \geq 1 - \delta,$$

when $|S| \geq \frac{18}{\epsilon^2}(\ln |\mathcal{F}'| + \ln \frac{2}{\delta})$. Plugging in $|\mathcal{F}'| = \left\lceil \frac{12\pi Kn^{c+1}}{\epsilon} \right\rceil^{O(n^{c+1})}$, we can prove that the hypothesis space \mathcal{F} is PAC learnable with sample complexity

$$\nu(\epsilon, \delta) = O \left(\frac{1}{\epsilon^2} \left(n^{c+1} \ln \frac{n}{\epsilon} + \ln \frac{1}{\delta} \right) \right).$$

This completes the proof of theorem 2. \square

We denote \mathcal{E} as the collection of all the functions $f_C : \mathcal{X} \rightarrow \mathcal{Y}$, where $C \in \mathcal{C}$. In fact, we can prove that \mathcal{E} is PAC learnable as well. By theorem 1, our hypothesis space \mathcal{F} can cover all the quantum circuits in \mathcal{C} . Thus we can obtain that $\mathcal{E} \subseteq \mathcal{F}$. Using the inequality (6), we will arrive at

$$\mathbb{P}_S \left(\forall f_h \in \mathcal{E} : |\hat{R}(f_h) - R(f_h)| \leq \frac{\epsilon}{2} \right) \geq 1 - \delta, \quad (7)$$

when $|S| \geq \frac{18}{\epsilon^2} (\ln |\mathcal{F}'| + \ln \frac{2}{\delta})$. Similarly with the method in theorem 2, we combine the inequality (7) with proposition 6, and we can prove that \mathcal{E} is PAC learnable with sample complexity $\nu(\epsilon, \delta) = O(\frac{1}{\epsilon^2} (n^{c+1} \ln \frac{n}{\epsilon} + \ln \frac{1}{\delta}))$ as well.

We stress the differences between our results and the previous works [28, 33] in the literature. First, in reference [28] Chung and Lin focused on a *finite* set of discretized quantum channels, and their algorithm is based on random orthogonal measurements. Whereas, in our settings we focus on a set of unitary quantum circuits with continuous variational parameters, thus the size of our concept class is *infinite*. Moreover, our proof is based on a family of variational quantum neural networks with an explicit uniform structure, which would be useful in practical applications. Second, in reference [33] Bu *et al* considered a more general class of quantum channels and their bounds of the sample complexity grow exponentially with the number of qubits n . In contrast, our focus here is variational quantum circuits and the sample complexity we obtained scales only polynomially with the system size. In other words, while reference [33]'s setting is more general, the sample complexity bounds obtained in this work is exponentially tighter. Our work and references [28, 33] are complementary to each other.

It is also worthwhile to clarify that, although we have proved that the sample complexity for learning any physical quantum circuit is low (namely, it only scales polynomially with the number of qubits involved), this does not mean that these circuits can be learned efficiently since the time complexity to learn an unknown circuit can still be exponentially high. In fact, it has been proved recently that training a variational quantum circuit, even for logarithmically many qubits and free fermionic systems, is NP-hard [43]. This implies that although we know for sure that our hypothesis \mathcal{F} can cover all physical quantum circuits and only a polynomial number of samples are needed to train a variational circuit $F \in \mathcal{F}$, how to efficiently solve the optimization problem of minimizing the empirical risk remains unclear and might be an exponentially hard problem in practice.

3. Discussion

We mention that the family of hypothesis quantum variational circuits constructed in this paper is of independent interest due to its use of only $O(n^{c+1})$ variational parameters while maintaining notable representation power. These circuits might be used as variational ansatz for implementing quantum classifiers [23, 25, 44–48], variational quantum eigensolvers [49–52], or quantum generative adversarial networks [53–55], etc. On the other hand, we also remark that similarly to many other variational quantum circuits constructed in the literature, this family of variational circuits may suffer from the barren plateau (i.e., vanishing gradient) problem [56, 57] as well. In addition, our work can be appealing as the family of circuits is constructed without optimizing the structure and the number of parameters. In the future, it would also be interesting to explore other alternative structures with smaller depths and fewer parameters. Another interesting problem worth further investigation is to consider a scenario where we do not have perfect knowledge about the training data, namely that the training dataset may not be fully labeled. How to extend our results to this scenario remains unknown.

We note that in our proof, the use of PAC learning theory is in fact independent from the learning model, i.e., it can deal with both the classical and quantum objectives. In our settings, the objects to be learned are parametric quantum circuits, but we can still use standard classical techniques of PAC learning theory (like discretization) to obtain the sample complexity bound.

In summary, we have proved that unitary physical quantum circuits are PAC learnable on a quantum computer via empirical risk minimization. In particular, we proved that to learn a unitary quantum circuit with at most n^c local gates, the sample complexity is bounded by $\tilde{O}(n^{c+1})$. Our results are generally applicable to all unitary quantum circuits of practical interest. There are many notable quantum circuits (algorithms or kernels, such as Shor's factorization algorithm [39], the HHL algorithm [12], quantum support vector machine [14], quantum classification based on discrete logarithm [58], etc) that hold the intriguing potential of exponential quantum speedup. Our results imply that a polynomial number of samples are enough to learn these quantum circuits. In reference [59], Bang *et al* proposed a method for learning quantum algorithms assisted by machine learning, which shows learning speedup in designing

quantum circuits for solving the Deutsch–Jozsa problem, and our results imply that the quantum circuits they used are PAC learnable as well.

Acknowledgments

We thank Wenjie Jiang, Peixin Shen, and Xun Gao in particular for their helpful discussions. This work is supported by the start-up fund from Tsinghua University (Grant No. 53330300320), the National Natural Science Foundation of China (Grant No. 12075128), and the Shanghai Qi Zhi Institute.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Appendix A. The universality of \mathcal{F}

In this paper, all the constants such as b , c , K , and M are independent of n , ϵ , and δ . Also, we recall that \mathcal{C} is the set of all the n -qubit quantum circuits with at most n^c unitary gates, with each gate acting on at most b qubits.

In proving theorem 1 in the main text, we used three lemmas, which are appended in the following. The lemma 1 is proved in reference [40], which we recap here for completeness. The propositions 1 and 2 are proved in this paper.

Lemma 1 ([40], section 4.5.2). *An arbitrary unitary operation on b qubits can be implemented using a circuit containing at most $c_0 b^2 4^b$ single-qubit unitary gates and CNOT gates, where c_0 is a constant.*

Proposition 1. *For any $C \in \mathcal{C}$, there exist $l = O(n^c)$ unitary gates U_1, U_2, \dots, U_l , such that $C = U_l U_{l-1} \dots U_1$, and each gate U_i is either a single-qubit unitary gate V_j acting on qubit j , or CNOT_{1j} gate with the first qubit being the control qubit and the j th qubit being the target one.*

Proof. We first prove that C can be decomposed as $O(n^c)$ elementary gates, including CNOT gates and single-qubit unitary gates. By lemma 1, C can be implemented by at most $c_0 b^2 4^b n^c = O(n^c)$ unitary gates, and each gate is either a single-qubit unitary gate or CNOT_{ij} gate with the control qubit i and the target qubit j .

To prove that C can be decomposed as the product of $l = O(n^c)$ single-qubit unitary gates and CNOT_{1j} gates, we need only prove that when $i \neq 1$ and $i \neq j$, CNOT_{ij} can be decomposed as CNOT_{1i} , CNOT_{1j} , and H gates.

When $i > 1, j = 1$, we can write CNOT_{i1} in this way:

$$\text{CNOT}_{i1} = (H_1 \otimes H_i) \text{CNOT}_{1i} (H_1 \otimes H_i).$$

Meanwhile, when $i, j > 1$ and $i \neq j$, we can decompose CNOT_{ij} into CNOT_{1j} and CNOT_{i1} in this way:

$$\text{CNOT}_{ij} = (\text{CNOT}_{1j} \text{CNOT}_{i1})^2,$$

and we have shown that CNOT_{i1} can be decomposed as CNOT_{1i} and H gates.

As each decomposition uses only $O(1)$ gates, we can obtain that C can be decomposed as the product of $l = O(n^c)$ single-qubit unitary gates and CNOT_{1j} gates, and the proof is completed. \square

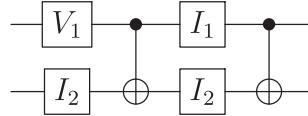
In a level-2 block B_j , there are two level-1 blocks on qubit 1 and qubit j , respectively. Each level-1 block has three parameters β, γ, δ , and by Z – X decomposition [40], we can tune these three parameters to enable a level-1 block L_j on qubit j to represent any single-qubit unitary gate acting on qubit j up to an irrelevant global phase. Also, by setting the three parameters to zero, a level-1 block L_j can also represent the identity gate. We will prove that by tuning the parameters of the level-2 blocks, one layer consisting of $B_n B_{n-1} \dots B_2$ can represent any single-qubit unitary gate acting on any qubit j , and CNOT_{1j} can be represented by two layers.

Proposition 2. (1) *One layer $B_n B_{n-1} \dots B_2$ can represent any single-qubit unitary gate V_j acting on any qubit j up to an irrelevant global phase.*

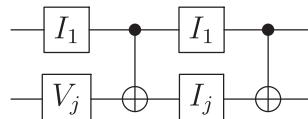
(2) *$(B_n B_{n-1} \dots B_2)^2$ can represent CNOT_{1j} up to an irrelevant global phase.*

Proof. To prove this lemma, we will set most of the level-2 blocks in the layers to be the identity gates and use at most two blocks to represent the gates we need.

We prove part one first. We separate the claim into two cases, $j = 1$ and $j \neq 1$. When $j = 1$, we can let $B_n B_{n-1} \dots B_3$ represent the identity gate by tuning all their parameters to zero. For clarity, we denote L_j as a level-1 block acting on the j th qubit. Given any unitary gate V_1 on qubit 1, a level-1 block L_1 can represent V_1 , and both level-1 blocks L_1 and L_2 can represent the identity gate. As a level-2 block B_2 consists of four level-1 blocks and two CNOT gates, we can tune the parameters of the four level-1 blocks in the following way so that B_2 can represent V_1 :



Similarly, when $j \neq 1$, we can let $B_n B_{n-1} \dots B_{j+1}$ and $B_{j-1} \dots B_3 B_2$ represent the identity gate. Then we need only let B_j represent unitary gate V_j . Given any unitary gate V_j on qubit j , we can tune the parameters of the four level-1 blocks in the following way so that B_j can represent V_j :

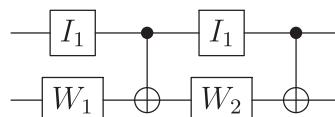


Therefore, the proof of part one is completed. Now we will prove part two. We set all the parameters in the two layers $(B_n B_{n-1} \dots B_2)^2$ to be zero except the two B_j blocks. Then we will use two B_j blocks to represent CNOT_{1j} . We decompose CNOT_{1j} up to an irrelevant global phase factor $e^{-i\pi/4}$ in the following way:

$$(W_4 \otimes W_3) \text{CNOT}_{1j} (I_1 \otimes W_2) \text{CNOT}_{1j} (I_1 \otimes W_1),$$

where we set $W_1 = R_z\left(\frac{\pi}{2}\right)$, $W_2 = R_y\left(\frac{\pi}{2}\right)$, $W_3 = R_z\left(-\frac{\pi}{2}\right) R_y\left(-\frac{\pi}{2}\right)$, and $W_4 = R_z\left(-\frac{\pi}{2}\right)$. Here we denote $R_z(\theta) = e^{-i\theta Z/2}$ and $R_y(\theta) = e^{-i\theta Y/2}$ as the rotation operators along the z -axis and y -axis on the Bloch sphere, respectively. In addition, W_4 and the identity gate I_1 act on the first qubit, and W_1 , W_2 , and W_3 act on the j th qubit.

Hence, we use the B_j block in the first layer to represent $\text{CNOT}_{1j} (I_1 \otimes W_2) \text{CNOT}_{1j} (I_1 \otimes W_1)$ in this way:



Finally, we use the second level-2 block B_j to represent $W_4 \otimes W_3$, where W_4 acts on the first qubit and W_3 acts on the j th qubit. Therefore, two layers $(B_n B_{n-1} \dots B_2)^2$ can represent CNOT_{1j} up to an irrelevant global phase, and this completes the proof of part two. \square

Appendix B. PAC learnability of \mathcal{F}

The following lemma shows that any finite hypothesis space \mathcal{F}' is PAC learnable.

Lemma 2 ([42], corollary 1.2). *Assume that the hypothesis space \mathcal{F}' is finite, $\delta \in (0, 1]$, $\epsilon > 0$ and the range of the loss function is in an interval of length $c \geq 0$. Then if the size of the training set $|S| \geq \frac{c^2}{2\epsilon^2} (\ln |\mathcal{F}'| + \ln \frac{2}{\delta})$, the event $\forall f_h \in \mathcal{F}' : |\hat{R}(f_h) - R(f_h)| \leq \epsilon$ holds with probability at least $1 - \delta$ over repeated sampling of the training set S .*

Our circuit F consists of $R_x(\theta)$, H , and CNOT gates. By assigning two sets of different values to the variational parameters θ , we can get two distinct circuits F_1 and F_2 , and their corresponding hypothesis functions f_1 and f_2 are different. We note that although F_1 and F_2 differ in the value of their variational parameters θ , their ordering of the gates ($R_x(\theta)$, H , and CNOT gates) are the same. We will show that when all the variational parameters in circuit F_1 and F_2 are close enough, the risk and empirical risk of f_1 and f_2 will be close. To prove this, we first define the distance of two unitary matrices $U_1, U_2 \in \mathbb{C}^{2^n \times 2^n}$ as the 2-norm of the matrix $U_1 - U_2$:

$$E(U_1, U_2) = \|U_1 - U_2\|_2 = \sup_{x \in \mathcal{X}} \|(U_1 - U_2)|\psi(x)\rangle\|_2.$$

Now we introduce the following lemma about the function $E(U_1, U_2)$.

Proposition 3. *The function $E(U_1, U_2)$ satisfies the following properties:*

- (a) *Let U, V be the $R_x(\theta), R_x(\theta + \epsilon)$ gates acting on the j th qubit, respectively, where $\epsilon \in (0, 1), \theta \in [0, 2\pi], j = 1, 2, \dots, n$. Then $E(U, V) \leq \epsilon$.*
- (b) *$E(U_1 U_{l-1} \dots U_1, V_l V_{l-1} \dots V_1) \leq \sum_{j=1}^l E(U_j, V_j)$, where $U_1, U_2, \dots, U_l, V_1, V_2, \dots, V_l$ are unitary matrices.*

Proof. The second property is shown in [40], section 4.5.3. We need only prove the first property.

$$\begin{aligned}
 E(U, V) &= \|U - V\|_2 \\
 &= \|R_x(\theta) - R_x(\theta + \epsilon)\|_2 \\
 &= \|I - R_x(\epsilon)\|_2 \\
 &\stackrel{(i)}{=} \|I - (I - i\epsilon X/2 + (i\epsilon X/2)^2/(2!) - \dots)\|_2 \\
 &\leq \|i\epsilon X/2\|_2 + \|(i\epsilon X/2)^2/(2!)\|_2 + \dots \\
 &\stackrel{(ii)}{=} \exp(\epsilon/2) - 1 \\
 &\leq \epsilon,
 \end{aligned}$$

where (i) uses Taylor's expansion of the operator $R_x(\epsilon) = e^{-i\epsilon X/2}$, and (ii) uses Taylor's expansion of $\exp(\epsilon/2)$ and that $\|X\|_2 = \|I\|_2 = 1$.

We recall that $\mathcal{L}(y_1, y_2)$ is the trace distance of two pure states $|\psi(y_1)\rangle$ and $|\psi(y_2)\rangle$. Then we introduce the following properties of $\mathcal{L}(y_1, y_2)$.

Proposition 4. *The function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, 1]$ satisfies the following two properties:*

- (a) *For any $y_1, y_2, y_3 \in \mathcal{Y}$, we have $\mathcal{L}(y_1, y_3) - \mathcal{L}(y_2, y_3) \leq \mathcal{L}(y_1, y_2)$.*
- (b) *For any $y_1, y_2 \in \mathcal{Y}$, we have $\mathcal{L}(y_1, y_2) \leq \|\psi(y_1)\rangle - |\psi(y_2)\rangle\|_2$.*

Proof. The first part of this lemma is the triangle inequality, which is proved in [40], section 9.2.1. Here, we only prove the second property. We denote $F(|\psi(y_1)\rangle, |\psi(y_2)\rangle) = |\langle\psi(y_1)|\psi(y_2)\rangle|$ as the fidelity between the two states $|\psi(y_1)\rangle$ and $|\psi(y_2)\rangle$. Then we will arrive at

$$\begin{aligned}
 \mathcal{L}(y_1, y_2) &= \frac{1}{2} \|\psi(y_1)\rangle - |\psi(y_2)\rangle\|_1 \\
 &\stackrel{(iii)}{=} \sqrt{1 - F(|\psi(y_1)\rangle, |\psi(y_2)\rangle)^2},
 \end{aligned}$$

where the proof of equation (iii) is given in [40], section 9.2.3.

In addition, we note that for any complex number $z \in \mathbb{C}$ and its complex conjugate $z^* \in \mathbb{C}$, as $(|z| - 1)^2 \geq 0$, we have $2 - 2|z| \geq 1 - |z|^2$. Hence, we get $2 - z - z^* \geq 2 - 2|z| \geq 1 - |z|^2$. Let $z = \langle\psi(y_1)|\psi(y_2)\rangle$, we obtain that

$$\begin{aligned}
 \mathcal{L}(y_1, y_2) &= \sqrt{1 - F(|\psi(y_1)\rangle, |\psi(y_2)\rangle)^2} = \sqrt{1 - |z|^2} \\
 &\leq \sqrt{2 - z - z^*} = \|\psi(y_1)\rangle - |\psi(y_2)\rangle\|_2.
 \end{aligned}$$

This completes the proof. □

Now we will use the properties of $E(U_1, U_2)$ and $\mathcal{L}(y_1, y_2)$ to show that the differences of both the risk and empirical risk between f_1 and f_2 are bounded by ϵ , where the hypothesis functions f_1 and f_2 correspond to the variational circuits F_1 and F_2 , respectively.

Proposition 5. *We denote $\boldsymbol{\theta}^{F_1} = (\theta_1^{F_1}, \theta_2^{F_1}, \dots, \theta_l^{F_1})$ as a vector containing all the variational parameters in F_1 , where l is the number of $R_x(\theta)$ gates in circuit F , and $\theta_i^{F_1}$ is the value of the variational parameter characterizing the i th x-rotation of F_1 . Similarly, we denote $\boldsymbol{\theta}^{F_2} = (\theta_1^{F_2}, \theta_2^{F_2}, \dots, \theta_l^{F_2})$ as a vector containing all the variational parameters in F_2 .*

Let $f_1, f_2 \in \mathcal{F}$ be the corresponding hypothesis functions of F_1, F_2 , respectively. Then given any probability distribution P over $\mathcal{X} \times \mathcal{Y}$ and training set $S = (x_i, y_i)_{i=1,2,\dots,m}$, the following two inequalities hold if $\|\boldsymbol{\theta}^{F_1} - \boldsymbol{\theta}^{F_2}\|_\infty \leq \frac{\epsilon}{Kn^{c+1}}$ (K is a large enough constant):

$$\begin{aligned}
 |R(f_1) - R(f_2)| &\leq \epsilon, \\
 |\hat{R}(f_1) - \hat{R}(f_2)| &\leq \epsilon.
 \end{aligned}$$

Proof. First, we will prove that $E(F_1, F_2) \leq \epsilon$ when $\|\theta^{F_1} - \theta^{F_2}\|_\infty \leq \frac{\epsilon}{Kn^{c+1}}$. Then we will use it to show the risk and empirical risk of f_1 and f_2 are close.

As F is composed of H gates, $R_x(\theta)$ gates and CNOT gates, we can write $F_1 = U_l U_{l-1} \dots U_1$ and $F_2 = V_l V_{l-1} \dots V_1$, where U_i is the i th gate in F_1 , and V_i is the i th gate in F_2 . As U_i and V_i are of the same type of gates, we can prove that $E(U_i, V_i) \leq \frac{\epsilon}{Kn^{c+1}}$ by separating different cases on the types of U_i and V_i :

Case I: If U_i and V_i are both H gates or both CNOT gates, as there is no variational parameter in H or CNOT, we have $U_i = V_i$, and we obtain that $E(U_i, V_i) = 0$.

Case II: If U_i and V_i are both $R_x(\theta)$ gates, as the difference of $\theta_i^{F_1}$ and $\theta_i^{F_2}$ is at most $\frac{\epsilon}{Kn^{c+1}}$, by the first property of proposition 3, we have $E(U_i, V_i) \leq \frac{\epsilon}{Kn^{c+1}}$.

We note that $l = O(n^{c+1})$ by our construction of F . By the second property of proposition 3 and choosing a large enough constant K such that $l \leq Kn^{c+1}$, we can get

$$\begin{aligned} E(F_1, F_2) &= E(U_l U_{l-1} \dots U_1, V_l V_{l-1} \dots V_1) \\ &\leq \sum_{j=1}^l E(U_j, V_j) \\ &\leq \sum_{j=1}^l \frac{\epsilon}{Kn^{c+1}} \leq \epsilon. \end{aligned}$$

Now we can bound the differences of the risk and empirical risk between the two hypothesis functions f_1 and f_2 , respectively. For convenience, we define $D(f_1, f_2)$ as $\sup_{x \in \mathcal{X}, y \in \mathcal{Y}} |\mathcal{L}(y, f_1(x)) - \mathcal{L}(y, f_2(x))|$. We observe that both $|R(f_1) - R(f_2)|$ and $|\hat{R}(f_1) - \hat{R}(f_2)|$ can be bounded by $D(f_1, f_2)$. Hence, we will prove that $D(f_1, f_2) \leq \epsilon$, and we can obtain the two inequalities $|R(f_1) - R(f_2)| \leq \epsilon$ and $|\hat{R}(f_1) - \hat{R}(f_2)| \leq \epsilon$.

$$\begin{aligned} D(f_1, f_2) &\stackrel{(iv)}{\leq} \sup_{x \in \mathcal{X}} \mathcal{L}(f_1(x), f_2(x)) \\ &\stackrel{(v)}{\leq} \sup_{x \in \mathcal{X}} \|f_1(x) - f_2(x)\|_2 \\ &= \sup_{x \in \mathcal{X}} \|(F_1 - F_2)|\psi(x)\rangle\|_2 \\ &= E(F_1, F_2) \leq \epsilon, \end{aligned}$$

where (iv) uses the first property of function \mathcal{L} in proposition 4, and (v) uses the second property of function \mathcal{L} in proposition 4. This completes the proof of proposition 5. \square

We note that in our proof of theorem 2, we used lemma 2 and proposition 5 to show that $\forall f_h \in \mathcal{F} : |\hat{R}(f_h) - R(f_h)| \leq \frac{\epsilon}{2}$ holds with probability $1 - \delta$. To prove that \mathcal{F} is PAC learnable, we introduce the following technical lemma.

Proposition 6. Assume $\forall f_h \in \mathcal{F} : |\hat{R}(f_h) - R(f_h)| \leq \frac{\epsilon}{2}$ holds. We denote $h = \arg \min_{h' \in \mathcal{F}} R(h')$, and $\hat{h} = \arg \min_{h' \in \mathcal{F}} \hat{R}(h')$. Then we have

$$R(\hat{h}) - R(h) \leq 2 \sup_{h' \in \mathcal{F}} |\hat{R}(h') - R(h')| \leq \epsilon.$$

Proof. The proof of this inequality is given in [42], section 1.2. We give the proof of the lemma here for completeness. To bound $R(\hat{h}) - R(h)$, we observe that it can be expressed as the sum of $R(\hat{h}) - \hat{R}(\hat{h})$ and $\hat{R}(\hat{h}) - R(h)$. Then we can use $\sup_{h' \in \mathcal{F}} |\hat{R}(h') - R(h')|$ to bound $R(\hat{h}) - \hat{R}(\hat{h})$ and $\hat{R}(\hat{h}) - R(h)$, respectively.

$$\begin{aligned}
R(\hat{h}) - R(h) &= R(\hat{h}) - \hat{R}(\hat{h}) + \hat{R}(\hat{h}) - R(h) \\
&\stackrel{(vi)}{=} R(\hat{h}) - \hat{R}(\hat{h}) + \sup_{h' \in \mathcal{F}} (\hat{R}(\hat{h}) - R(h')) \\
&\stackrel{(vii)}{\leqslant} R(\hat{h}) - \hat{R}(\hat{h}) + \sup_{h' \in \mathcal{F}} (\hat{R}(h') - R(h')) \\
&\leqslant |\hat{R}(\hat{h}) - R(\hat{h})| + \sup_{h' \in \mathcal{F}} |\hat{R}(h') - R(h')| \\
&\leqslant 2 \sup_{h' \in \mathcal{F}} |\hat{R}(h') - R(h')| \leqslant \epsilon,
\end{aligned}$$

where (vi) uses that $h = \arg \min_{h' \in \mathcal{F}} R(h')$, and (vii) uses that $\hat{h} = \arg \min_{h' \in \mathcal{F}} \hat{R}(h')$. \square

ORCID iDs

Haoyuan Cai  <https://orcid.org/0000-0003-1315-1361>

Dong-Ling Deng  <https://orcid.org/0000-0002-1042-4646>

References

- [1] Jordan M and Mitchell T 2015 *Science* **349** 255–60
- [2] LeCun Y, Bengio Y and Hinton G 2015 *Nature* **521** 436–44
- [3] Silver D *et al* 2016 *Nature* **529** 484–9
- [4] Silver D *et al* 2017 *Nature* **550** 354–9
- [5] Senior A W *et al* 2020 *Nature* **577** 706–10
- [6] Krizhevsky A, Sutskever I and Hinton G E 2017 *Commun. ACM* **60** 84–90
- [7] Sarma S D, Deng D L and Duan L M 2019 *Phys. Today* **72** 48
- [8] Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N and Lloyd S 2017 *Nature* **549** 195
- [9] Ciliberto C, Herbster M, Davide Lalongo A, Pontil M, Rocchetto A, Severini S and Wossnig L 2017 *Proc. R. Soc. A* **474** 20170551
- [10] Dunjko V and Briegel H J 2018 *Rep. Prog. Phys.* **81** 074001
- [11] Carleo G, Cirac J, Cranmer K, Daudet L, Schuld M, Tishby N, Vogt-Maranto L and Zdeborová L 2019 *Rev. Mod. Phys.* **91** 045002
- [12] Harrow A W, Hassidim A and Lloyd S 2009 *Phys. Rev. Lett.* **103** 150502
- [13] Gao X, Zhang Z Y and Duan L M 2018 *Sci. Adv.* **4** eaat9004
- [14] Rebentrost P, Mohseni M and Lloyd S 2014 *Phys. Rev. Lett.* **113** 130503
- [15] Aaronson S 2015 *Nat. Phys.* **11** 291–3
- [16] Valiant L G 1984 *Commun. ACM* **27** 1134–42
- [17] Haussler D 1990 Probably Approximately Correct Learning *The Eighth National Conf. on Artificial Intelligence Boston July, 1990*
- [18] Shalev-Shwartz S and Ben-David S 2014 *Understanding Machine Learning: From Theory to Algorithms* (Cambridge: Cambridge University Press)
- [19] Arute F *et al* 2019 *Nature* **574** 505–10
- [20] Zhong H S *et al* 2020 *Science* **370** 1460–3
- [21] Wu Y *et al* 2021 Strong quantum computational advantage using a superconducting quantum processor (arXiv:2106.14734)
- [22] Schuld M, Fingerhuth M and Petruccione F 2017 *Europhys. Lett.* **119** 60002
- [23] Schuld M, Bocharov A, Svore K M and Wiebe N 2020 *Phys. Rev. A* **101** 032308
- [24] Beer K, Bondarenko D, Farrelly T, Osborne T J, Salzmann R, Scheiermann D and Wolf R 2020 *Nat. Commun.* **11** 808
- [25] Cong I, Choi S and Lukin M D 2019 *Nat. Phys.* **15** 1273–8
- [26] Watts A B, Kothari R, Schaeffer L and Tal A 2019 Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits *Proc. 51st Annual ACM SIGACT Symp. on Theory of Computing STOC 2019* (New York: ACM) pp 515–26
- [27] Arunachalam S and de Wolf R 2017 Optimal quantum sample complexity of learning algorithms *Proc. 32nd Computational Complexity Conf. CCC '17* (Dagstuhl, DEU: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik)
- [28] Chung K M and Lin H H 2021 Sample efficient algorithms for learning quantum channels in PAC model and the approximate state discrimination problem *16th Conf. on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)* (Leibniz International Proceedings in Informatics (LIPIcs) vol 197) ed M H Hsieh (Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik) pp 3:1–3:22 <https://drops.dagstuhl.de/opus/volltexte/2021/13998>
- [29] Arunachalam S and de Wolf R 2017 *SIGACT News* **48** 41–67
- [30] Arunachalam S, Grilo A B and Yuen H 2020 Quantum statistical query learning (arXiv:2002.08240)
- [31] Heidari M, Padakandla A and Szpankowski W 2021 A theoretical framework for learning from quantum data (arXiv:2107.06406)
- [32] Sweke R, Seifert J P, Hangleiter D and Eisert J 2021 *Quantum* **5** 417
- [33] Bu K, Koh D E, Li L, Luo Q and Zhang Y 2021 On the statistical complexity of quantum circuits (arXiv:2101.06154)
- [34] Caro M C and Datta I 2020 *Quantum Mach. Intell.* **2** 1–14
- [35] Du Y, Tu Z, Yuan X and Tao D 2021 An efficient measure for the expressivity of variational quantum algorithms (arXiv:2104.09961)
- [36] Caro M C, Gil-Fuster E, Meyer J J, Eisert J and Sweke R 2021 Encoding-dependent generalization bounds for parametrized quantum circuits (arXiv:2106.03880)
- [37] Cheng H C, Hsieh M H and Yeh P C 2015 arXiv:1501.00559
- [38] Chitambar E and Gour G 2019 *Rev. Mod. Phys.* **91** 025001

- [39] Shor P 1994 Algorithms for quantum computation: discrete logarithms and factoring *Proc. 35th Annual Symp. on Foundations of Computer Science* pp 124–34
- [40] Nielsen M A and Chuang I L 2010 *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge: Cambridge University Press)
- [41] Coppersmith D 2002 An approximate Fourier transform useful in quantum factoring (arXiv:[quant-ph/0201067](https://arxiv.org/abs/quant-ph/0201067))
- [42] Wolf M M 2020 Mathematical foundations of supervised learning https://m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MA4801_2020S/ML_notes_main.pdf
- [43] Bittel L and Kliesch M 2021 Training variational quantum algorithms is np-hard-even for logarithmically many qubits and free fermionic systems (arXiv:[2101.07267](https://arxiv.org/abs/2101.07267))
- [44] Farhi E and Neven H 2018 Classification with quantum neural networks on near term processors (arXiv:[1802.06002](https://arxiv.org/abs/1802.06002))
- [45] Havlíček V, Córcoles A D, Temme K, Harrow A W, Kandala A, Chow J M and Gambetta J M 2019 *Nature* **567** 209
- [46] Zhu D *et al* 2019 *Sci. Adv.* **5** eaaw9918
- [47] Grant E, Benedetti M, Cao S, Hallam A, Lockhart J, Stojovic V, Green A G and Severini S 2018 *npj Quantum Inf.* **4** 65
- [48] Lu S, Duan L M and Deng D L 2020 *Phys. Rev. Res.* **2** 033212
- [49] Peruzzo A, McClean J, Shadbolt P, Yung M H, Zhou X Q, Love P J, Aspuru-Guzik A and O'Brien J L 2014 *Nat. Commun.* **5** 4213
- [50] Kokail C *et al* 2019 *Nature* **569** 355
- [51] Liu J G, Zhang Y H, Wan Y and Wang L 2019 *Phys. Rev. Res.* **1** 023025
- [52] Wang D, Higgott O and Brierley S 2019 *Phys. Rev. Lett.* **122** 140504
- [53] Lloyd S and Weedbrook C 2018 *Phys. Rev. Lett.* **121** 040502
- [54] Dallaire-Demers P L and Killoran N 2018 *Phys. Rev. A* **98** 012324
- [55] Hu L *et al* 2019 *Sci. Adv.* **5** eaav2761
- [56] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 *Nat. Commun.* **9** 4812
- [57] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 *Nat. Commun.* **12** 1791
- [58] Liu Y, Arunachalam S and Temme K 2021 *Nat. Phys.* **17** 1013–7
- [59] Bang J, Ryu J, Yoo S, Pawłowski M and Lee J 2014 *New J. Phys.* **16** 073017