

# New Bounds on Quantum Sample Complexity of Measurement Classes

Mohsen Heidari  
Department of Computer Science  
Indiana University  
Email: mheidar@iu.edu

Wojciech Szpankowski  
Department Computer Science  
Purdue University  
Email: szpan@purdue.edu

**Abstract**—This paper studies quantum supervised learning for classical inference from quantum states. In this model, a learner has access to a set of labeled quantum samples as the training set. The objective is to find a quantum measurement that predicts the label of the unseen samples. The hardness of learning is measured via sample complexity under a quantum counterpart of the well-known *probably approximately correct* (PAC). Quantum sample complexity is expected to be higher than classical one, because of the measurement incompatibility and state collapse. Recent efforts showed that the sample complexity of learning a finite quantum concept class  $\mathcal{C}$  scales as  $O(|\mathcal{C}|)$ . This is significantly higher than the classical sample complexity that grows logarithmically with the class size. This work improves the sample complexity bound to  $O(V_{\mathcal{C}^*} \log |\mathcal{C}^*|)$ , where  $\mathcal{C}^*$  is the set of extreme points of the convex closure of  $\mathcal{C}$  and  $V_{\mathcal{C}^*}$  is the *shadow-norm* of this set. We show the tightness of our bound for the class of bounded Hilbert–Schmidt norm, scaling as  $O(\log |\mathcal{C}^*|)$ . Our approach is based on a new quantum *empirical risk minimization* (ERM) algorithm equipped with a shadow tomography method.

## I. INTRODUCTION

Quantum learning is one of the leading applications of quantum computing both for classical and quantum problems. While some models suggest quantum-enhancements of classical learning by mapping data into input quantum states [1]–[4], quantum computers (QCs) have a far greater capability to learn patterns from inherently quantum data. This is possible by directly operating on quantum states of physical systems (e.g., photons or states of matter) or their qubit representations [5]–[12]. Learning from quantum data has been studied extensively in recent literature in the context of diverse applications, including quantum simulation [10]–[14], phase-of-matter detection [5], [15], ground-state search [6], [7], [16], and entanglement detection [8], [9], [17]–[20]. In this context, variational quantum algorithms (VQAs) and quantum neural networks (QNNs) are among the most conventional models of a quantum learner that consist of a parameterized circuit trained in a hybrid quantum-classical loop.

One of the objectives of quantum learning is to train a model with a minimum number of iterations or samples. This is crucial in scenarios where quantum data is scarce or is expensive to produce. For instance, sometimes in the quantum simulation of ground state search, the quantum states are the output of a Hamiltonian with high gate complexity, making

the state preparation expensive. An explicit example is the Hamiltonian of the two-dimensional Fermi-Hubbard model on an  $8 \times 8$  lattice that requires about  $10^7$  Toffoli gates [21]. The challenge is exacerbated as the training models are data-intensive due to the *no-cloning* and *state collapse* that prohibit sample reuse — a phenomenon unique to quantum. Therefore, studying the quantum sample complexity, as a measure of the hardness of training a model, is foundational to our understanding of quantum learning. This is the focus of our study.

In classical learning theory, sample complexity has been studied extensively for decades under the well-known PAC framework [22], [23]. The theory of quantum learning has been studied recently under various models [24]–[30]. A relevant example is *quantum state discrimination*, where the objective is to distinguish an unknown quantum state  $\rho$  from another (known or unknown) state using *measurements* on multiple samples [25], [27], [31]–[36]. In the Bshouty and Jackson [26], a classical PAC problem is studied using a quantum oracle that outputs identical copies of an associated superposition state [26], [37]–[40].

In this work, we consider *quantum probably approximately correct* (QPAC) which is a natural generalization of PAC in quantum [29]. QPAC applies to conventional quantum supervised learning with VQAs, where one aims to accurately predict a classical property (a.k.a., label) of an unknown quantum state. QPAC is a generic framework that subsumes several models such as state discrimination, quantum property testing, quantum state classification, and classical PAC. A learner has access to  $n$  physical quantum states with their classical label as the supervised training set. The learner’s objective is to find a quantum measurement that minimizes the loss for predicting the label of unseen samples. Here, the learner does not know the sample distribution, the classical description of the states, and the labeling law. It can only measure the  $n$  physical samples. With the *no-free lunch* theorem, the learner can only hope to find a predictor with a loss sufficiently close to the best candidate in a predefined set called the *concept class*. QNNs or variational circuits are examples of concept classes. The quantum sample complexity of a given concept class is the minimum sample size  $n$  such that a learner successfully finds a predictor performing close to the best candidate of the class.

QPAC is therefore a distribution-free and state-free requirement. It is stronger than PAC, as PAC is only distribution-free. QPAC abides by quantum mechanical laws such as no-cloning, state collapse, and measurement incompatibility. Such properties prohibit sample reuse and, thus, raise new challenges for learning in quantum settings. Moreover, quantum models (i.e., QNNs) are significantly richer than classical models. Therefore, given the fragility of quantum samples, the strictness of QPAC, and the richness of quantum models, one expects quantum sample complexity to be significantly, if not exponentially, greater than the classical one.

Recent results show that any finite concept class  $\mathcal{C}$  is learnable with  $\mathcal{O}(|\mathcal{C}|)$  quantum samples [29] leaving an exponential gap compared to classical sample complexity that scales with  $\mathcal{O}(\log |\mathcal{C}|)$ . This gap is mainly due to the no-cloning and state collapse of quantum. However, a better bound is remained to be shown. Infinite classes have been studied recently in [30], [41]. In [41] the results of [42] was extended to infinite concept classes through an  $\epsilon$ -netting argument. The authors in [30] studied learning of  $k$ -qubit (or  $k$ -local) operators i.e., operations acting non-trivially on at most  $k$  out of  $d$  qubits. They showed that the sample complexity scales as  $\mathcal{O}(k 4^k \log d)$ . When  $k \ll d$ , this bound is comparable to  $\mathcal{O}(k \log d)$  which is the sample complexity of classical  $k$ -juntas.

#### A. Main Contributions

We study learning of generic finite or infinite quantum concept classes. Our main contributions are two-fold.

**Bounds on quantum sample complexity.** This paper presents a new improved bound on the quantum sample complexity of generic concept classes (finite or infinite). Particularly, we prove that the quantum sample complexity of any measurement concept class  $\mathcal{C}$  grows with  $\mathcal{O}(V_{\mathcal{C}}^* \log |\mathcal{C}^*|)$ , where  $\mathcal{C}^*$  is the set of extreme points of the convex closure of the concept class and  $V_{\mathcal{C}}^*$  is the *shadow-norm* of this set. In general,  $|\mathcal{C}^*| \leq |\mathcal{C}|$ . In an extreme case  $|\mathcal{C}|$  could be infinite while  $|\mathcal{C}^*|$  is finite. The bound is tight and simplifies to  $\mathcal{O}(\log |\mathcal{C}|)$  for concept classed with bounded Hilbert–Schmidt norm. These results are surprising especially since sample duplication is prohibited and measurement incompatibility would lead to an exponentially larger sample complexity with standard methods. Such results are steps toward a quantum analogous to the fundamental theorem of statistical learning, which remains open.

**QSRM algorithm.** We propose a new approach called *quantum shadow risk minimization* (QSRM) to measure the empirical risk of the predictors in the class that substantially improves the quantum sample complexity of existing works. Our approach relies on searching only the *extreme* points of the concept class and making use of shadow tomography [43]. In classical learning ERM is a brute-force search to minimize the empirical loss and has a sample complexity of  $\mathcal{O}(\log |\mathcal{C}|)$ . Extending this algorithm to quantum is not straightforward. One could naively propose the same technique to compute the empirical risk of each quantum predictor and choose the

one with the minimum risk. However, the no-cloning and measurement incompatibility makes this approach prohibitive. Essentially, the training samples will be distorted each time we measure the empirical risk of a predictor. If we use that naively, then we might need fresh samples for each predictor, leading to a sample complicity of  $\mathcal{O}(|\mathcal{C}|)$ . There have been multiple attempts [29], [41] to introduce an efficient quantum ERM, obtaining a better bound that grows with  $\mathcal{O}(\log |\mathcal{C}|)$  when measurements in  $\mathcal{C}$  are fully *compatible* and  $\mathcal{O}(|\mathcal{C}|)$  in fully *incompatible* class. Our algorithm achieves a sample complexity  $\mathcal{O}(V_{\mathcal{C}}^* \log |\mathcal{C}^*|)$  that improves upon these works.

## II. MODEL FORMULATION

**Notations:** For shorthand, denote  $[d]$  as  $\{1, 2, \dots, d\}$ . As usual, a quantum measurement  $\mathcal{M}$  is a positive operator-valued measure (POVM) represented by a set of operators  $\mathcal{M} := \{M_v, v \in \mathcal{V}\}$ , where  $\mathcal{V}$  is the set of possible outcomes,  $M_v \geq 0$  for any  $v \in \mathcal{V}$ , and  $\sum_{v \in \mathcal{V}} M_v = I$ . Throughout of this paper, we assume a finite-dimensional Hilbert space  $\mathcal{H}$  with dimensionality denoted by  $\dim(\mathcal{H})$ .

### A. Quantum Learning Model

Our model follows that of [29]. The objective is to distinguish between multiple groups of unknown quantum states without prior knowledge about the states. Available is only a training set of quantum states with a classical label determining their group index. We seek a procedure that learns the labeling law. The model is defined more precisely as follows.

Let  $\mathcal{Y}$  denote the labeling set and  $\mathcal{H}$  be the underlying finite-dimensional Hilbert space. There are  $n$  sample  $(|\phi\rangle_i, y_i), i \in [n]$  randomly generated according to an unknown but fixed probability distribution  $D$ <sup>1</sup>. A predictor is a quantum measurement  $\mathcal{M} := \{M_{\hat{y}} : \hat{y} \in \mathcal{Y}\}$  that acts on the quantum states and outputs  $\hat{y} \in \mathcal{Y}$  as the predicted label. From Born’s rule,  $\hat{y}_i$  is generated randomly with probability  $\langle \phi_i | M_{\hat{y}_i} | \phi_i \rangle$ . The prediction loss is determined via a loss function  $l : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, \infty)$ . Therefore, the generalization (expected) loss is calculated as  $L_D(\mathcal{M}) = \mathbb{E}[l(Y, \hat{Y})]$ , where the expectation is taken over all the randomness involved.

The learner’s objective is to minimize  $L_D$  while having no prior knowledge about  $D$  and the structure of the states. It can only measure the  $n$  physical samples without knowing their classical descriptions. Due to the *no-free lunch* theorem, the learner can only hope to find a  $\mathcal{M}$  with  $L_D(\mathcal{M})$  sufficiently close to the best candidate within a fixed collection  $\mathcal{C}$  of predictors called the *concept class*.

**Definition 1 (QPAC).** A quantum learning algorithm agnostically QPAC learns a concept class  $\mathcal{C}$  if there exists a function  $n_{\mathcal{C}} : (0, 1)^2 \mapsto \mathbb{N}$  such that for every  $\epsilon, \delta \in [0, 1]$  and given  $n > n_{\mathcal{C}}(\epsilon, \delta)$  samples drawn independent and identically distributed (i.i.d.) according to any probability distributions  $D$  and any unknown states  $|\phi\rangle_i, i \in [n]$ , the algorithm outputs, with probability of at least  $(1 - \delta)$ , a measurement whose

<sup>1</sup>Alternatively, one can view the samples as copies of the mixed state  $\rho_{XY} := \mathbb{E}_D[|\phi\rangle\langle\phi| \otimes |Y\rangle\langle Y|]$ .

loss is less than  $\inf_{\mathcal{M} \in \mathcal{C}} L_D(\mathcal{M}) + \epsilon$ .<sup>2</sup> The quantum sample complexity of  $\mathcal{C}$  is the minimum of  $n_{\mathcal{C}}$  for which there exists a QPAC learner.

The state discrimination problem is a special case in which samples are identical and are either of two *a priori* known states. QPAC also subsumes classical PAC as classical samples and concept classes are mutually diagonalized in the canonical basis. Generally, QPAC is a stronger requirement than PAC and other methods, as it is a *agnostic*, distribution-free and state-free condition. Hence, one expects that quantum sample complexity is significantly higher than classical one.

### B. Related Works on QERM

It is known that classical ERM PAC learns any (classical) finite concept class  $\mathcal{C}$  with sample complexity that scales with  $\mathcal{O}\left(\frac{1}{\epsilon^2} \log \frac{|\mathcal{C}|}{\delta}\right)$ . In quantum settings, there have been various attempts in developing counterparts of ERM algorithm. Due to the no-cloning theorem, the straightforward quantum extension of ERM results in a sample complexity of  $\mathcal{O}\left(\frac{|\mathcal{C}|}{\epsilon^2} \log \frac{1}{\delta}\right)$ , see [29] for more details. This is problematic as the sample complexity grows linearly with the size of the concept class.

In [29], a new ERM-type algorithm is introduced to improve this bound. The new bound depends on the measurement *incompatibility* structure of the concept class. Incompatible measurements cannot be measured simultaneously (for more details see [44]). On one extreme, all the measurements in the concept class are mutually compatible; on another extreme, there is no pair of compatible measurements. Based on this, an improved bound on sample complexity is as follows.

**Fact 1** ([29]). *Quantum sample complexity of any finite concept class  $\mathcal{C}$  is upper bounded as*

$$n_{\mathcal{C}}(\epsilon, \delta) \leq \min_{\mathcal{C}_r, \text{Comp. partition}} \sum_{r=1}^b \left\lceil \frac{8}{\epsilon^2} \log \frac{2b|\mathcal{C}_r|}{\delta} \right\rceil,$$

where the minimization is taken over all compatibility partitioning of  $\mathcal{C}$  and  $b$  is the number of partition bins. This bound ranges from  $\mathcal{O}\left(\frac{1}{\epsilon^2} \log \frac{|\mathcal{C}|}{\delta}\right)$ , for fully compatible class, to  $\mathcal{O}\left(\frac{|\mathcal{C}|}{\epsilon^2} \log \frac{1}{\delta}\right)$  for fully incompatible class.

In [41], this result was extended to infinite concept classes through an  $\epsilon$ -netting argument. In this paper, we propose a new quantum ERM that substantially improves the above bounds to one that, in the worst case, grows with  $\mathcal{O}\left(\frac{1}{\epsilon^2} \log \frac{|\mathcal{C}|}{\delta}\right)$  even for fully incompatible concept classes.

## III. MAIN RESULTS

In this section, we present the main results of the paper on the sample complexity. We prove the upper bound by proposing a quantum shadow risk minimization (QSRM). We first discuss methods to measure the empirical loss.

<sup>2</sup>Naturally, we are interested in efficient learning where  $n_{\mathcal{C}}$  is polynomial in  $\epsilon$ ,  $\delta$  and  $\dim(\mathcal{H})$ .

### A. Measuring the Empirical Loss

Without loss of generality, assume  $l : \mathcal{Y} \times \mathcal{Y} \mapsto [0, 1]$ . Let  $\mathcal{Z}$  be the image set of  $l$ , that is  $\mathcal{Z} := \{l(y, \hat{y}) : y, \hat{y} \in \mathcal{Y}\}$ . Since  $\mathcal{Y}$  is a finite set, then so is  $\mathcal{Z}$ . We define a loss observable for any predictor  $\mathcal{M}$  as  $\mathcal{L}_{\mathcal{M}} := \{L_z^{\mathcal{M}} : z \in \mathcal{Z}\}$ , where

$$L_z^{\mathcal{M}} = \sum_{y, \hat{y} \in \mathcal{Y}} \mathbb{1}_{\{l(y, \hat{y})=z\}} M_{\hat{y}} \otimes |y\rangle\langle y|, \quad \forall z \in \mathcal{Z}. \quad (1)$$

With that, the expectation value of  $\mathcal{L}_{\mathcal{M}}$  is equal to  $L_D(\mathcal{M})$ . Moreover, applying  $\mathcal{L}_{\mathcal{M}}$  on any sample gives the incurred empirical loss of  $\mathcal{M}$  on that sample.

Now given a set of samples  $\mathcal{S}_n = \{(|\phi_i\rangle, y_i) : i \in [n]\}$ , we would like to estimate the empirical loss for all predictors in  $\mathcal{C}$ . In the classical setting, this can be done easily by testing each predictor on the training samples, leading to a  $\mathcal{O}(\log |\mathcal{C}|)$  bound on  $n$ . However, in quantum, the samples can be used once as they collapse after the measurements. Moreover, the loss measurements  $\mathcal{L}_{\mathcal{M}}$  might be *incompatible* for different  $\mathcal{M} \in \mathcal{C}$  — rendering simultaneous measuring obsolete. As a result, a naive ERM in quantum would need fresh copies for each predictor, needing  $\mathcal{O}(|\mathcal{C}|)$  samples.

In what follows, we introduce a new approach for measuring the empirical loss. Our approach is inspired by shadow tomography [43] that is applied to *identical copies* of quantum states. Shadow tomography is a procedure to obtain useful classical information via random measuring of the input state. In what follows we explain this procedure.

Consider a generic state  $\rho$  in a Hilbert space  $\mathcal{H}$ . Generate a unitary operator  $U$  randomly from a class of choices  $\mathcal{U}$  to be determined. Apply  $U$  on the input state resulting in the state  $U^\dagger \rho U$ . Next, measure the rotated state in the canonical basis  $|j\rangle, j \in [\dim(\mathcal{H})]$ . From Born's rule the probability of getting the output  $j$  is  $p_j = \langle j|U^\dagger \rho U|j\rangle$ . Given an outcome  $j$ , prepare the state  $\omega_j = U |j\rangle\langle j| U^\dagger$ . The expectation  $\mathbb{E}_{\sim(\mathcal{J}, U)}[\omega_j]$  over the measurement randomness ( $p_j$ ) and the choice of  $U$  equals to  $\Gamma[\rho]$ , where  $\Gamma$  is a mapping defined as

$$\Gamma[O] := \mathbb{E}_U \left[ \sum_{j \in [\dim(\mathcal{H})]} \langle j|U^\dagger O U|j\rangle U |j\rangle\langle j| U^\dagger \right], \quad (2)$$

for any operator  $O$  on  $\mathcal{H}$ . Observe that  $\Gamma$  is a linear mapping on  $\mathcal{B}(\mathcal{H})$  and hence has an inverse denoted by  $\Gamma^{-1}$ . We note that  $\Gamma^{-1}$  is the shadow channel  $\mathcal{M}^{-1}$  introduced in [43]. We apply  $\Gamma^{-1}$  on  $\omega_j$  resulting in the so called shadow

$$\hat{\rho} := \Gamma^{-1}[U |j\rangle\langle j| U^\dagger]. \quad (3)$$

Note that  $\hat{\rho}$  is a classical matrix and hence can be copied several times. Moreover,  $\hat{\rho}$  is not a valid density operator as it is not necessarily a positive semi-definite matrix. However, it is an unbiased estimate of the original state.

**Fact 2** ([43]). *When  $\mathcal{U}$  is tomographically complete, the classical shadow  $\hat{\rho}$  is unbiased, that is  $\mathbb{E}_{U, j}[\hat{\rho}] = \rho$ .*

For our problem, we apply the above process with a random unitary  $U_i$  for each sample  $\rho_i := |\phi_i\rangle\langle\phi_i|$  resulting in  $n$




Fig. 1. The process for estimating the empirical loss of each measurement.

classical shadows  $\hat{\rho}_i, i \in [n]$ . Next, we compute the expected loss per shadow as

$$\hat{L}_i(M) = \sum_{\hat{y}} l(y, \hat{y}) \text{tr}\{M_{\hat{y}} \hat{\rho}_i\}.$$

This process is demonstrated in Figure 1. Next, we average it over all the shadows to find

$$L_{\hat{S}_n}(\mathcal{M}) := \frac{1}{n} \sum_{i=1}^n \hat{L}_i(M), \quad (4)$$

where  $\hat{S}_n$  represents the set of all the shadows  $\hat{\rho}_i$  with their labels  $y_i$ . Since the classical shadows can be copied, this procedure estimates the empirical loss of all the predictors in the concept class without asking for fresh samples as in the naive approach.

**Remark 1.** The estimate  $L_{\hat{S}_n}$  is unbiased for any predictor  $\mathcal{M}$ , that is  $\mathbb{E}[L_{\hat{S}_n}] = L_D(\mathcal{M})$ .

**Remark 2.** Huang et al. [43] used the median of means estimator for predicting the expectation of arbitrary observables. In our work, we slightly deviate and consider the empirical mean estimator to make sure our sample complexity bounds have expressions comparable with standard PAC bounds.

## B. Concentration bounds

We prove that a similar result as in [43] also holds for the empirical average estimator, as opposed to the median of means. We proceed with presenting the following definition.

**Definition 2.** The shadow norm of any operator  $O$  on  $\mathcal{H}$  is

$$\|O\|_{\text{shadow}} := \max_{\sigma \in \mathcal{D}[\mathcal{H}]} \left( \sum_{j \in [\dim(\mathcal{H})]} \langle j|U^\dagger \sigma U|j\rangle \langle j|U\Gamma^{-1}[O]U^\dagger|j\rangle^2 \right)^{1/2}.$$

Depending on the choice of  $\mathcal{U}$  different bounds are obtained for shadow norm. When  $\mathcal{U}$  is the Clifford group, Shadow norm is related to the Hilbert-Schmidt norm [43]:

$$\|O\|_{\text{shadow}} \leq \sqrt{3 \text{tr}\{O^2\}}. \quad (5)$$

When we use random Pauli measurements, the shadow norm relates to the locality of an operator. More precisely, for an operator  $O$  that acts non-trivially on only  $k$  out of  $d$  qubits, the shadow norm is bounded as  $\|O\|_{\text{shadow}} \leq 2^k \|O\|_\infty$ . Our first result extends Theorem 1 of [43] to empirical mean estimators.

**Proposition 1.** Let  $\hat{\rho}_i, i \in [n]$  be the classical shadows of  $n$  copies of a mixed state  $\rho$ , as in (3). Then, given  $\delta \in (0, 1)$  and  $m$  arbitrary observables  $O_1, \dots, O_m$ , the empirical average  $\hat{\delta}_j := \frac{1}{n} \sum_i \text{tr}\{O_j \hat{\rho}_i\}$  satisfies the additive error  $\varepsilon$  given that

$$n \geq \mathcal{O} \left( \frac{1}{\varepsilon^2} \log \left( \frac{m}{\delta} \right) \max_i \left\| O_i - \frac{\text{tr}\{O_i\}}{\dim(\mathcal{H})} I \right\|_{\text{shadow}}^2 \right).$$

The proof is given in Section IV-A. The proposition implies that the median-of-means can be replaced with an empirical averaging without performance loss. Similar observations have been made regarding 3-design circuits and the entire unitary group for classical shadow tomography [45], [46]. Recall that the objective of QSRM is to find  $\min_{\mathcal{M}} L_{\hat{S}_n}$ . However, the upper bound in the above theorem becomes loose when the concept class is large, for example when  $|\mathcal{C}|$  is infinite. Next, We show that only the extreme points of  $\mathcal{C}$  are relevant.

## C. Extreme points of a concept class

Let  $\bar{\mathcal{C}}$  denote the convex closure (envelope) of  $\mathcal{C}$ . Note that  $\bar{\mathcal{C}}$  is the set of all POVMs that can be written as a convex combination of measurements in  $\mathcal{C}$ . More precisely, POVMs of the form  $\bar{\mathcal{M}} = \{\bar{M}_{\hat{y}}, \hat{y} \in \mathcal{Y}\}$  such that

$$\bar{M}_{\hat{y}} = \sum_{j=1}^k \alpha_j M_{\hat{y}}^j, \quad \forall \hat{y} \in \mathcal{Y},$$

where each  $M^j = \{M_{\hat{y}}^j, \hat{y} \in \mathcal{Y}\}$  belongs to  $\mathcal{C}$ , and  $\alpha_j \in [0, 1]$  with  $\sum_j \alpha_j = 1$ . By definition,  $\bar{\mathcal{C}}$  is a convex set. Therefore, it must have extreme points.

**Definition 3.** Given any concept class  $\mathcal{C}$ , by  $\mathcal{C}^*$  denote the extreme points of its convex closure.

With this perspective, QSRM is a brute-force search inside  $\mathcal{C}^*$  and is equipped with the shadow empirical mean estimator of (4). This is summarized as Algorithm 1 followed by Theorem 1 proving its QPAC learnability.

---

### Algorithm 1: QSRM

---

**Input:**  $\mathcal{C}^*$  of the concept class and  $n$  training samples.

- 1 **for**  $i = 1$  **to**  $n$  **do**
  - 2     Generate a unitary  $U_i$  randomly.
  - 3     Apply  $U_i$  on  $\rho_i$  as in Figure 1.
  - 4     Measure along  $\{|j\rangle, j \in [\dim(\mathcal{H})]\}$  to get  $j_i$ .
  - 5     Store the classical shadows  $\hat{\rho}_i$  as in (3).
  - 6 **for each**  $\mathcal{M}$  in  $\mathcal{C}^*$  **do**
  - 7     Compute the empirical loss  $L_{\hat{S}_n}(\mathcal{M})$  as in (4).
  - 8 **return**  $\hat{\mathcal{M}}$  with the minimum  $L_{\hat{S}_n}$ .
- 

**Theorem 1.** Suppose  $\ell$  is a bounded loss function and  $\mathcal{C}$  is a measurement class with finite extreme points. Then, QSRM (Algorithm 1) agnostically QPAC learns  $\mathcal{C}$  with quantum sample complexity bounded as

$$n_{\mathcal{C}}(\varepsilon, \delta) = \mathcal{O} \left( \frac{V_{\mathcal{C}^*}}{\varepsilon^2} \log \frac{|\mathcal{C}^*|}{\delta} \right),$$



where  $\mathcal{C}^*$  is the set of extreme points as in Definition 3 and

$$V_{\mathcal{C}^*} := \max_{\mathcal{M} \in \mathcal{C}^*} \max_y \left\| \mathcal{L}_{\mathcal{M}}(y) - \frac{\text{tr}\{\mathcal{L}_{\mathcal{M}}(y)\}}{\dim(\mathcal{H})} I \right\|_{\text{shadow}}^2,$$

is the shadow norm as in Definition 2.

Note that  $|\mathcal{C}^*| \leq |\mathcal{C}|$ . Moreover,  $\mathcal{C}^*$  can be finite even when  $\mathcal{C}$  is infinite. This leads to an interesting distinction compared to the classical learning: the quantum sample complexity of  $\mathcal{C}$  can scale with a rate smaller than  $\log |\mathcal{C}|$ . It is worth noting that even though the Hilbert space is finite-dimensional,  $\mathcal{C}^*$  could be infinite. In that case, one can combine QSRM with an  $\epsilon$ -netting procedure to get a bound on the sample complexity.

The theorem implies a tight sample complexity bound, up to a constant, for classes of bounded norm.

**Corollary 1.** *The sample complexity of any concept class  $\mathcal{C}$ , such that  $|\mathcal{C}^*|$  is finite and the Hilbert–Schmidt norm of each  $\mathcal{M} \in \mathcal{C}$  is bounded by an independent constant, scales as  $\Theta(\log |\mathcal{C}^*|)$ .*

*Proof.* The upper bounds follows from Theorem 1 and (5). The lower bound is implied from the fact that quantum sample complexity is not smaller than classical one, and that the classical sample complexity of a finite concept class  $\mathcal{C}_{\text{classical}}$  scales as  $\Theta(\log |\mathcal{C}_{\text{classical}}|)$ .  $\square$

#### IV. PROOF OF THE MAIN RESULTS

##### A. Proof of Proposition 1

The proof follows from a concentration of measures for bounded variances [47, Theorem 8.2]:

**Lemma 1** (Method of Bounded Variances). *Let  $X_1, \dots, X_n$  be a set of random variables and let  $Y_n = f(X_1, \dots, X_n)$ , where  $f$  a function such that  $\mathbb{E}[Y_n] < \infty$ . Let  $D_i = \mathbb{E}[Y_n | \mathbf{X}^i] - \mathbb{E}[Y_n | \mathbf{X}^{i-1}]$  and  $|D_i| \leq c_i$  for some constants  $c_i > 0$ . Also let,  $V_n = \sum_{i=1}^n \sup_{\mathbf{x}^{i-1}} \text{var}(D_i | \mathbf{x}^{i-1})$ . Then,*

$$\mathbb{P}\left\{ |Y_n - \mathbb{E}[Y_n]| > \epsilon \right\} \leq 2 \exp\left\{ -\frac{\epsilon^2}{4V_n} \right\},$$

where  $\epsilon \leq \frac{2V_n}{\max_i c_i}$ .

**Corollary 2.** *For any  $O$ , and small enough  $\epsilon > 0$ , the shadow empirical loss in (4) satisfies*

$$\mathbb{P}\left\{ |\hat{\delta} - \text{tr}\{O\rho\}| > \epsilon \right\} \leq 2 \exp\left\{ \frac{-n\epsilon^2}{4 \max_i \text{var}(\text{tr}\{O\hat{\rho}_i\})} \right\}.$$

*Proof.* The corollary is proved from Lemma 1 with  $f = \hat{\delta} := \frac{1}{n} \sum_i \text{tr}\{O\hat{\rho}_i\}$  implying that  $D_i = \frac{1}{n}(\text{tr}\{O\hat{\rho}_i\} - \mathbb{E}[\text{tr}\{O\hat{\rho}_i\}])$ . Since,  $\hat{\rho}_i$ 's are mutually independent and identically distributed then  $V_n = \frac{1}{n^2} \sum_i \text{var}(\text{tr}\{O\hat{\rho}_i\}) \leq \frac{1}{n} \max_i \text{var}(\text{tr}\{O\hat{\rho}_i\})$ , which gives the desired statement.  $\square$

Applying this result for  $O = O_i$  in the proposition and a union bound give  $\max_j |\hat{\delta}_j - \text{tr}\{O_j \rho\}| \leq \epsilon$  with probability greater than

$$1 - 2m \exp\left\{ \frac{-n\epsilon^2}{4 \max_{i,j} \text{var}(\text{tr}\{O_j \hat{\rho}_i\})} \right\}. \quad (6)$$

From [43, Lemma 1], the variance terms are upper bounded by the shadow norm as  $\max_j \left\| O_j - \frac{\text{tr}\{O_j\}}{\dim(\mathcal{H})} I \right\|_{\text{shadow}}^2$ . Equating the bound on the probability to  $\delta$  gives the desired result.

##### B. Proof of Theorem 1.

We start with the following lemma:

**Lemma 2.** *Let  $\bar{\mathcal{C}}$  be the convex closure of  $\mathcal{C}$  and  $\mathcal{C}^*$  be the set of all extreme points of  $\bar{\mathcal{C}}$ . Then,  $\text{opt}_{\mathcal{C}} = \text{opt}_{\bar{\mathcal{C}}} = \inf_{\mathcal{M} \in \mathcal{C}^*} L_D(\mathcal{M})$ .*

*Proof.* Note that  $L_D(\mathcal{M})$  is linear in  $\mathcal{M}$  as it is equal to the expectation of the loss observable  $\mathcal{L}_{\mathcal{M}}$ . This is due to the linearity of the trace and the definition of  $\mathcal{L}_{\mathcal{M}}$ . As a result, given that  $\mathcal{C} \subseteq \bar{\mathcal{C}}$  and that  $\text{opt}_{\bar{\mathcal{C}}} = \inf_{\mathcal{M} \in \bar{\mathcal{C}}} L_D(\mathcal{M})$  we find that  $\text{opt}_{\bar{\mathcal{C}}} = \text{opt}_{\mathcal{C}}$ . Moreover, since the above expression is a convex optimization, then the optimal values occur at the extreme points of  $\bar{\mathcal{C}}$ . Hence the proof is complete.  $\square$

This result implies that QPAC learning of  $\mathcal{C}$  is reduced to its extreme points  $\mathcal{C}^*$ . The proof of Theorem 1 follows from Proposition 1 and Lemma 2 with the following argument. For any fixed  $y$ , and  $\mathcal{M}$  let  $\mathcal{L}_{\mathcal{M}}(y)$  be the expected loss when the label is  $y$ . We apply Proposition 1 with  $O_j = \mathcal{L}_{\mathcal{M}}(y)$  for  $\mathcal{M} \in \mathcal{C}^*$ . Let

$$V_{\mathcal{C}^*} := \max_{\mathcal{M} \in \mathcal{C}^*} \max_y \left\| \mathcal{L}_{\mathcal{M}}(y) - \frac{\text{tr}\{\mathcal{L}_{\mathcal{M}}(y)\}}{\dim(\mathcal{H})} I \right\|_{\text{shadow}}^2.$$

Then, the proposition gives the following sample complexity bound:  $n = \mathcal{O}\left(\frac{V_{\mathcal{C}^*}}{\epsilon^2} \log\left(\frac{|\mathcal{C}^*|}{\delta}\right)\right)$ . Now, let  $\widehat{\mathcal{M}}$  and  $\mathcal{M}^*$  be the measurements minimizing  $L_{\widehat{\mathcal{S}}_n}$  and  $L_D$ , respectively. Then, with probability  $(1 - \delta)$  we have that:

$$L_D(\widehat{\mathcal{M}}) \leq \widehat{L}(\widehat{\mathcal{M}}) + \frac{\epsilon}{2} \leq \widehat{L}(\mathcal{M}^*) + \frac{\epsilon}{2} \leq L_D(\mathcal{M}^*) + \epsilon.$$

The left-hand side is the loss of the selected predictor by QSRM (Algorithm 1), and the right-hand side equals  $\text{opt}_{\mathcal{C}} + \epsilon$  and hence the proof is complete.

#### CONCLUSION

This paper studies the learning of quantum measurement classes. It introduces a novel quantum algorithm called QSRM for learning quantum concept classes. Using this algorithm, a new upper bound on the quantum sample complexity is derived. It is shown that the quantum sample complexity grows at most with the logarithm of the size of the extreme points of the convex closure of the concept class. This is a significant improvement over prior results. The approach is based on a novel method to estimate the empirical loss of the concept class via creating random shadows of the training samples. With that QSRM algorithm can perform risk minimization while abiding to no-cloning, state collapse, and measurement incompatibility.

#### ACKNOWLEDGMENT

This work was partially supported by the NSF Center for Science of Information (CSoI) Grant CCF-0939370, and in addition by NSF Grants CCF-1524312, CCF-2006440, and CCF-2211423.

## REFERENCES

- [1] V. Giovannetti, S. Lloyd, and L. Maccone, “Quantum random access memory,” *Physical Review Letters*, vol. 100, no. 16, apr 2008.
- [2] D. K. Park, F. Petruccione, and J.-K. K. Rhee, “Circuit-based quantum random access memory for classical data,” *Scientific Reports*, vol. 9, no. 1, mar 2019.
- [3] S. Lloyd, M. Mohseni, and P. Rebentrost, “Quantum algorithms for supervised and unsupervised machine learning,” *arXiv:1307.0411*, 2013.
- [4] —, “Quantum principal component analysis,” vol. 10, no. 9, pp. 631–633, jul 2014.
- [5] J. Carrasquilla and R. G. Melko, “Machine learning phases of matter,” vol. 13, no. 5, pp. 431–434, feb 2017.
- [6] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” vol. 355, no. 6325, pp. 602–606, feb 2017.
- [7] M. Broughton, G. Verdon, T. McCourt, A. J. Martinez, J. H. Yoo, S. V. Isakov, P. Massey, R. Halavati, M. Y. Niu, A. Zlokapa, E. Peters, O. Lockwood, A. Skolik, S. Jerbi, V. Dunjko, M. Leib, M. Streif, D. V. Dollen, H. Chen, S. Cao, R. Wiersema, H.-Y. Huang, J. R. McClean, R. Babbush, S. Boixo, D. Bacon, A. K. Ho, H. Neven, and M. Mohseni, “Tensorflow quantum: A software framework for quantum machine learning,” *arXiv:2003.02989*, Mar. 2020.
- [8] F. V. Massoli, L. Vadicamo, G. Amato, and F. Falchi, “A leap among entanglement and neural networks: A quantum survey,” *arXiv:2107.03313*, Jul. 2021.
- [9] S. Lu, S. Huang, K. Li, J. Li, J. Chen, D. Lu, Z. Ji, Y. Shen, D. Zhou, and B. Zeng, “Separability-entanglement classifier via machine learning,” *Physical Review A*, vol. 98, no. 1, p. 012315, 2018.
- [10] I. Kassal, J. D. Whitfield, A. Perdomo-Ortiz, M.-H. Yung, and A. Aspuru-Guzik, “Simulating chemistry using quantum computers,” *Annual Review of Physical Chemistry*, vol. 62, no. 1, pp. 185–207, may 2011.
- [11] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry,” *Reviews of Modern Physics*, vol. 92, no. 1, p. 015003, mar 2020.
- [12] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, “Quantum chemistry in the age of quantum computing,” *Chemical Reviews*, vol. 119, no. 19, pp. 10856–10915, aug 2019.
- [13] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, “Quantum chemistry calculations on a trapped-ion quantum simulator,” *Physical Review X*, vol. 8, no. 3, p. 031022, jul 2018.
- [14] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, “Quantum algorithms for quantum chemistry and quantum materials science,” *Chemical Reviews*, vol. 120, no. 22, pp. 12685–12717, oct 2020.
- [15] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, “Machine learning quantum phases of matter beyond the fermion sign problem,” vol. 7, no. 1, aug 2017.
- [16] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” vol. 549, no. 7671, pp. 195–202, sep 2017.
- [17] Y.-C. Ma and M.-H. Yung, “Transforming bell’s inequalities into state classifiers with machine learning,” *npj Quantum Information*, vol. 4, no. 1, jul 2018.
- [18] B. C. Hiesmayr, “Free versus bound entanglement, a NP-hard problem tackled by machine learning,” *Scientific Reports*, vol. 11, no. 1, oct 2021.
- [19] C. Chen, C. Ren, H. Lin, and H. Lu, “Entanglement structure detection via machine learning,” *Quantum Science and Technology*, 2021.
- [20] D.-L. Deng, X. Li, and S. D. Sarma, “Quantum entanglement in neural network states,” *Physical Review X*, vol. 7, no. 2, p. 021021, 2017.
- [21] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, “Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization,” *Quantum*, vol. 4, p. 296, jul 2020.
- [22] M. J. Kearns, R. E. Schapire, and L. M. Sellie, “Toward efficient agnostic learning,” *Machine Learning*, vol. 17, no. 2-3, pp. 115–141, 1994.
- [23] L. G. Valiant, “A theory of the learnable,” *Communications of the ACM*, vol. 27, no. 11, pp. 1134–1142, nov 1984.
- [24] S. Aaronson, “The learnability of quantum states,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 463, no. 2088, pp. 3089–3114, sep 2007.
- [25] S. M. Barnett and S. Croke, “Quantum state discrimination,” *Advances in Optics and Photonics*, vol. 1, no. 2, p. 238, feb 2009.
- [26] N. H. Bshouty and J. C. Jackson, “Learning dnf over the uniform distribution using a quantum example oracle,” *SIAM Journal on Computing*, vol. 28, no. 3, pp. 1136–1153, 1998.
- [27] C. Badescu, R. O’Donnell, and J. Wright, “Quantum state certification,” in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. ACM, jun 2019.
- [28] A. Montanaro and R. de Wolf, “A survey of quantum property testing,” *Theory of Computing*, vol. 1, no. 1, pp. 1–81, 2016.
- [29] M. Heidari, A. Padakandla, and W. Szpankowski, “A theoretical framework for learning from quantum data,” in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- [30] M. Heidari and W. Szpankowski, “Learning k-qubit quantum operators via pauli decomposition,” in *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, F. Ruiz, J. Dy, and J.-W. van de Meent, Eds., vol. 206. PMLR, 25–27 Apr 2023.
- [31] R. O’Donnell and J. Wright, “Efficient quantum tomography,” in *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, jun 2016.
- [32] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, “Sample-optimal tomography of quantum states,” in *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, jun 2016.
- [33] S. Bubeck, S. Chen, and J. Li, “Entanglement is necessary for optimal quantum property testing,” in *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, nov 2020.
- [34] S. Gambs, “Quantum classification,” *0809.0444 [quant-ph]*, Sep. 2008.
- [35] M. Guta and W. Kotłowski, “Quantum learning: asymptotically optimal classification of qubit states,” *New Journal of Physics*, vol. 12, no. 12, p. 123032, dec 2010.
- [36] H.-C. Cheng, A. Winter, and N. Yu, “Discrimination of quantum states under locality constraints in the many-copy setting,” in *2021 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jul 2021.
- [37] S. Arunachalam and R. de Wolf, “A survey of quantum learning theory,” *arXiv:1701.06806*, 2017.
- [38] V. Kanade, A. Rocchetto, and S. Severini, “Learning dnfs under product distributions via  $\mu$ -biased quantum fourier sampling,” *arXiv:1802.05690v3*, 2019.
- [39] E. Bernstein and U. Vazirani, “Quantum complexity theory,” *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1411–1473, oct 1997.
- [40] R. A. Servedio and S. J. Gortler, “Equivalences and separations between quantum and classical learnability,” *SIAM J. Comput.*, vol. 33, no. 5, p. 1067–1092, May 2004.
- [41] A. Padakandla and A. Magner, “Pac learning of quantum measurement classes : Sample complexity bounds and universal consistency,” in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, G. Camps-Valls, F. J. R. Ruiz, and I. Valera, Eds., vol. 151. PMLR, 28–30 Mar 2022, pp. 11305–11319.
- [42] M. Heidari, T. A. Atif, and S. S. Pradhan, “Faithful simulation of distributed quantum measurements with applications in distributed rate-distortion theory,” in *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, jul 2019.
- [43] H.-Y. Huang, R. Kueng, and J. Preskill, “Predicting many properties of a quantum system from very few measurements,” *Nature Physics* 16, 1050–1057 (2020), Feb. 2020.
- [44] A. S. Holevo, *Quantum Systems, Channels, Information*. DE GRUYTER, jan 2012.
- [45] J. Helsen and M. Walter, “Thrifty shadow estimation: Reusing quantum circuits and bounding tails,” *Physical Review Letters*, vol. 131, no. 24, p. 240602, Dec. 2023.
- [46] A. Zhao, N. C. Rubin, and A. Miyake, “Fermionic partial tomography via classical shadows,” *Physical Review Letters*, vol. 127, no. 11, p. 110504, Sep. 2021.
- [47] D. P. Dubhashi and A. Panconesi, *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, Jun. 2009.