

Empirical Risk Minimization and Uniform Convergence for Probabilistically Observed and Quantum Measurement Hypothesis Classes

Abram Magner

Department of Computer Science
University at Albany, State University of New York
Albany, NY, USA
Email: amagner@albany.edu

Arun Padakandla

Eurecom, Nice, France
Email: padakand@eurecom.fr

We continue the study of the learnability of quantum measurement classes in the setting where the learner is given access only to prepared quantum states, aiming for necessary and sufficient conditions for PAC learnability, along with corresponding sample complexity bounds. In the quantum setting, in contrast with the classical probabilistically observed case, sampled states are perturbed when a quantum measurement is applied, according to the Born rule, so that distinct samples in the training data cannot be arbitrarily reused. We first probe the results from previous works on this setting. We show that the empirical risk defined in previous works and matching the definition in the classical theory can fail to satisfy the uniform convergence property enjoyed in the classical learning setting for classes that we can show to be PAC learnable. Moreover, we show that VC dimension generalization upper bounds in previous work are in many cases infinite, even for measurement classes defined on a finite-dimensional Hilbert space. We then show that, nonetheless, every measurement class defined on a finite-dimensional Hilbert space is PAC learnable via a modification of the ERM rule.

I. INTRODUCTION

Classical statistical learning theory formulates the broad problem of learning a relationship between two random quantities $X \in \mathcal{X}$ – known as features – and $Y \in \mathcal{Y}$ – class labels – as follows: the data are assumed to be generated from some *unknown* probability distribution $P_{X,Y}$, and a *learner* is given access to a dataset consisting of m independent and identically distributed samples (X_i, Y_i) . The learner’s task is to select a *hypothesis* from a fixed, known set Hyp of functions from \mathcal{X} to \mathcal{Y} (the *hypothesis/concept class*) that best approximates the joint distribution $P_{X,Y}$, *without knowledge of $P_{X,Y}$ itself*.

a) This work: The intent of this work is to make further progress in understanding learnability in the following supervised quantum learning scenario: there is an unknown joint probability distribution on prepared quantum states and classical labels. A hypothesis class consisting of quantum measurements is fixed and known to a learner. The learner is given access to a training dataset of these state-label pairs

sampled from the unknown data-generating distribution, but can only interact with the states by observing the classical outcomes of measuring them. It then outputs a hypothesis that is as close as possible to minimizing the expected *risk* over all hypotheses. This learning scenario was first posed in [1] as a quantum version of the classical PAC (Probably Approximately Correct) learning setting in which hypotheses are quantum measurements. This setting has extensive motivations ranging from building universal quantum state discriminators to classification of unknown quantum processes to classifying quantum phases of multipartite systems (see also [2], [3]). The setting was then further developed in [4]. In contrast with more well-established quantum learning frameworks [5], which deal with quantum algorithms for learning classical hypotheses (e.g., boolean functions $f : \{0,1\}^n \rightarrow \{0,1\}$) from superpositions of states corresponding to classical bit strings, our framework covers a distinct scenario in which input data consists of unknown quantum states, and the goal is to learn a measurement that predicts classical attributes (e.g., a class label) of those states.

More specifically, the authors of [1] formulated the quantum PAC learning framework that we study as follows: we fix a domain \mathcal{X} consisting of quantum states, along with a codomain \mathcal{Y} . Analytically, quantum states are described by density matrices¹ on a fixed Hilbert space \mathcal{H} over the complex numbers \mathbb{C} . We take the codomain $\mathcal{Y} = \{0,1\}$ for binary classification, but our results can be generalized further, for example to any finite codomain. A POVM hypothesis class Hyp is a set of *positive operator-valued measures* [6]², which specify quantum measurements with outcomes in \mathcal{Y} . Additionally, we fix a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, \infty]$. For binary classification, we take the misclassification loss $\ell(y_1, y_2) = I[y_1 \neq y_2]$, where $I[\cdot]$ is the indicator function. The learning process is as follows: an unknown distribution D on $\mathcal{X} \times \mathcal{Y}$ is fixed. To produce a single training example,

¹A density matrix is a positive semidefinite Hermitian matrix with trace 1.

²A 2-outcome POVM is a pair $\Pi := (\Pi_0, \Pi_1)$ of positive semidefinite operators on \mathcal{H} summing to the identity. Measurement of a state ρ with Π has outcome 0, 1 with probabilities $\text{Tr}(\rho\Pi_0)$, $\text{Tr}(\rho\Pi_1)$, respectively.

$(X, Y) \sim \mathcal{D}$ is sampled, and then a quantum register is prepared in state X . **Here and throughout, a quantum register is a collection of qubits prepared in a state that is a density matrix in \mathcal{H} , which may be multidimensional.** The learner is given access to the quantum register and Y , and can only interact with the quantum register by measuring it and observing the outcome. This occurs independently m times to produce a training set of size m . The learner is then allowed to make arbitrarily many measurements of the given quantum registers and by an arbitrary procedure then produces a resulting POVM h from the class Hyp. We note that each measurement alters the state of the register according to the axioms of quantum mechanics. The risk of a hypothesis is given by $R(h) = \mathbb{E}_{(X, Y) \sim \mathcal{D}}[\ell(h[X], Y)]$, where $h[X]$ denotes a random variable whose distribution is that of the outcome of measuring a quantum register in state X with POVM h . Then the goal of the learner is to output a hypothesis with risk close enough to the minimal risk achieved by any hypothesis in the class. We define this setup formally in Definition II.1 below.

The main problems of interest are similar to the ones asked in the classical PAC learning framework: perhaps the most immediate one is, what is a natural necessary and sufficient condition for PAC learnability of a POVM concept class? Is there a learning rule that is universal, in the sense that it is a PAC learning rule whenever the concept class is learnable? The present paper answers both of these questions. In the classical case with deterministic (function) concept classes, one of the fundamental results, which is sometimes called the *fundamental theorem of concept learning*, gives a necessary and sufficient condition for learnability of a concept class for binary classification under the misclassification loss: namely, learnability is equivalent to finiteness of the Vapnik-Chervonenkis (VC) dimension of the class [7]. The recent paper [4] gave one possible generalization of VC dimension to the quantum setting, resulting in a sufficient condition for learnability of POVM classes, along with a sample complexity³ upper bound for one particular learning rule. However, it gave no necessary conditions and did not explore the tightness of the upper bound or the universality of the learning rule. The present paper finds that this sufficient condition is substantially weak and that the learning rule is very far from universal.

A. Prior work

The literature on statistical problems involving quantum states and measurements is quite broad. For example, a wealth of quantum state estimation problems have been posed [8]–[10], wherein the input is a sequence of multiple quantum registers, all prepared in a single unknown state. This set of works also includes works on state tomography [11]–[19]. The task in such studies is to glean information about the single, unknown state – specifically, to *estimate* it. Estimation is *not* the same thing as learning, and so these are in contrast with our work, in which the goal is more analogous

³The *sample complexity* of a learning rule is the minimum number of samples required to guarantee that with probability at least $1 - \delta$, the risk (i.e., expected loss) of the learned hypothesis is within ϵ of the minimum possible.

to the classical supervised learning problem: i.e., our goal is to learn a statistical association between unknown quantum states sampled according to an unknown distribution and their classical labels. This statistical association need not reflect any intrinsic physical information about the states. We also point out that there are various works, such as [20], [21] that mix what is called PAC learning with quantum information, but these differ substantially from our setting: e.g., they assume a uniform distribution on the input, so they are not distribution-free; or they strongly constrain the input state to correspond to a bit string; or they output a boolean function instead of a POVM. There is also a large and expanding body of work in quantum machine learning in which hypothesis classes consist of specially structured POVMs – as a recent example, [22]. The focus in such works is different from that of the framework we study, since they aim to solve *classical* learning problems by suitably encoding classical input data as quantum states, then choosing a suitable measurement from the hypothesis class. In our case, the inputs \mathcal{X} are intrinsically quantum and are not encodings of known classical inputs.

At first glance, the paper [23] has a more related goal to ours – producing an optimal POVM from training data. However, training samples consist of the density matrices encoding states, rather than quantum registers, as well as the probabilities of outcomes of measurements by an unknown POVM. In contrast, in the framework that we consider, the inputs to a learner are not analytical state descriptions; rather, they are quantum registers prepared in those states. Furthermore, we are given, not probabilities of outcomes, but the outcomes themselves. Finally, the statistical relationship between the state and the label in our case can be arbitrary, whereas in the cited paper, it is governed by a single unknown POVM.

Two recent papers are the most relevant to the present one and, indeed, are the sources of the framework that we study in this paper: [1], [4]. The paper [1] formulated the POVM class PAC learning framework, showed that finite-cardinality POVM classes are PAC learnable, and pointed out the usefulness of joint measurability in reducing sample complexity, resulting in the *Quantum Empirical Risk Minimization (QERM)* learning rule. The QERM rule is a generalization of the classical ERM, which is the cornerstone of classical statistical learning theory.

The paper [4] studied the same setting, extending the sample complexity upper bounds for the QERM rule under the assumption that a partition is given, by formulating one possible generalization of the classical VC dimension of a probabilistically observed concept class. This implicitly showed that there exist PAC learnable POVM classes with infinite cardinality but left open the problem of giving necessary and sufficient conditions for a given class to be learnable. For example, no necessary conditions were given, in contrast with the present work. We will also show in this work that the upper bounds in that work are frequently vacuous.

In the course of proving our results, it will be convenient to define a PAC learning framework for what we call *probabilistically observed concept classes* (POCC), which we study as a technical tool for our quantum results. In this framework,

each concept is a function from \mathcal{X} to the set of probability distributions on \mathcal{Y} , and the task is, as usual to learn a risk-minimizing concept. However, on any sampled $x \in \mathcal{X}$ from the training set, for any concept h , the learner is only allowed to see a sample from the distribution $h(x)$. We will denote such samples by $h[x]$. This is in contrast with the theory of *probabilistic concepts* (p -concepts) introduced in [24]. There, concepts are similarly conditional distributions, but the learner is allowed to see the entire distribution $h(x)$.

B. Our contributions

- 1) **Results on failure of ERM and uniform convergence:** We first show that the natural ERM learning rule proposed and studied in [1], [4] can fail for probabilistically observed concept classes that are PAC learnable. We probe this phenomenon further, showing that the empirical risk can fail to satisfy the uniform convergence property for learnable hypothesis classes Hyp. That is, the supremal deviation of the empirical risk from expected value, where the supremum ranges over all elements of Hyp, does not converge to 0 as the number of samples tends to ∞ . This implies that in the probabilistically observed and the quantum case, the QERM learning rule cannot be universal in the sense of being a PAC learning rule if and only if the class to which it is applied is learnable.
- 2) **Learnability of finite dimensional hypothesis classes:** We then show that every POVM class defined on a finite-dimensional Hilbert space is PAC learnable. This implies that the nontrivial *qualitative* question of learnability/non-learnability only occurs in the infinite-dimensional case. Furthermore, this implies that recovering classical learning theory from the POVM class framework requires mapping of classical classes to POVM classes over infinite-dimensional Hilbert spaces. This is an indication that infinite-dimensional Hilbert spaces are of fundamental interest for a complete quantum learning theory.

Complete proofs, examples, further results, and extended discussions of other prior works are provided in <https://arxiv.org/abs/2308.12304>. Specifically, the proofs of our Theorems II.4, II.6, and III.3 are in Sections 7.1, 7.2, and 7.3, of the journal version, respectively. In Section 10, we give a more extended discussion of prior works and how they differ from ours, including, in particular, how works on channel tomography are not applicable to solve our learning problem.

II. MAIN RESULTS: LEARNABILITY, UNIFORM CONVERGENCE, AND ERM

A. Preliminaries

We first define the learning problems relevant to us. Definitions from quantum mechanics can be found in [6] and in the supplementary material.

Definition II.1 (POVM concept class learning problem [1], [4]). In the POVM concept/hypothesis class learning problem, we fix a set of possible input mixed states \mathcal{X} , which are density operators on a common Hilbert space \mathcal{H} , and a set of possible classical outputs \mathcal{Y} . We fix a *loss function* $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, \infty]$.

We fix a POVM concept class Hyp, which is simply a set of POVMs on \mathcal{H} having $|\mathcal{Y}|$ outcomes. Informally, a **learning rule** \mathcal{A} in this context takes as input a dataset $\{(\rho_j, Y_j)\}_{j=1}^m$ consisting of quantum registers in states $\rho_j \in \mathcal{X}$ and classical outputs $Y_j \in \mathcal{Y}$. This dataset is sampled from an unknown joint distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$. The learning rule interacts with the ρ_j via quantum measurements (formally, POVMs). Finally, it outputs a POVM $\Phi_* \in \text{Hyp}$ with the goal of minimizing $\mathbb{E}_{(X,Y) \sim \mathcal{D}}[\ell(\Phi_*[X], Y)]$, where $\Phi_*[X] \in \mathcal{Y}$ denotes the random outcome from measuring X with ρ_* .

We say that a POVM learning rule \mathcal{A} is (ϵ, δ) -probably approximately correct (PAC) for Hyp if there exists a sample size $m = m(\epsilon, \delta)$ such that for all distributions \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$ with $S \sim \mathcal{D}^m$, with probability at least $1 - \delta$, $\mathcal{A}(S)$ outputs a hypothesis $h \in \text{Hyp}$ satisfying $R(h) - \inf_{h_* \in \text{Hyp}} R(h_*) \leq \epsilon$.

We then say that Hyp is (ϵ, δ) -PAC learnable if there exists an (ϵ, δ) -PAC learning rule for Hyp. Finally, we say that Hyp is PAC learnable if it is (ϵ, δ) -PAC learnable for all $\epsilon, \delta > 0^4$.

The learning problem defined in Definition II.1 is related to the problem of *probabilistically observed concept class learning*, which we introduce below.

Definition II.2 (Probabilistically observed concept class learning problem). In the probabilistically observed concept class (POCC) learning problem, \mathcal{X} becomes an arbitrary set, and Hyp consists of functions $f : \mathcal{X} \rightarrow \Delta(\mathcal{Y})$, where $\Delta(S)$ denotes the set of probability distributions on a set S .

When a hypothesis $h \in \text{Hyp}$ is applied to an element $x \in \mathcal{X}$, the learning rule only observes a random sample $Z \sim h(x)$, not $h(x)$ itself. We denote a generic sample from $h(x)$ by $h[x]$.

Given this setting, the definition of PAC learning remains the same as before.

Remark II.3 (Probabilistic versus probabilistically observed concept learning). We emphasize the important distinction between the probabilistic concepts (also called p -concepts) of [25] and the probabilistically observed concepts in the present paper: in the p -concept framework, the output probability distribution itself is observed, rather than just a sample from it. In our setting, in contrast, our learning rules are only allowed to see a sample from an unknown output probability distribution.

1) *Connecting POVM classes with POCCs:* Here we describe the connection between the POVM and POCC frameworks. The POVM framework is more general than the POCC one: we first show how to translate the problem of learning

⁴The notion of PAC learnability formulated here is only concerned with learnability with *finite* sample complexity, not necessarily polynomial in the relevant variables. This is the same as in the classical definition of PAC learnability.

a POCC class to one of learning a POVM class, along with translations of POCC learning rules to POVM learning rules.

Given a POCC learning problem with domain \mathcal{X} and hypothesis class Hyp , the quantumization of this problem is formulated as follows: we introduce a Hilbert space \mathcal{H} with dimension equal to $|\mathcal{X}|$ (which may be uncountably infinite), and we choose, arbitrarily, an orthonormal basis $B = \{e_x\}_{x \in \mathcal{X}}$. Each $x \in \mathcal{X}$ corresponds to a basis element $e_x \in \mathcal{H}$. The domain of the POVM learning problem is the basis B . Each hypothesis $h \in \text{Hyp}$ bijectively maps to a corresponding POVM Π_h defined as follows: Π_h first measures in the basis B , uniquely identifying the input state e_x with probability 1, then postprocesses e_x through the classical channel corresponding to $h(x)$. (We note that a POVM may be constructed by measurement of a state with a POVM, then postprocessing the outcome through a classical channel.)

A POCC learning rule is a function of inputs x and samples from an arbitrary set of hypotheses $h[x]$. The analogous POVM learning rule is the same function as in the classical case, applied to the classical outcome of measurement in the basis B , along with results of passing this outcome through channels associated with hypotheses in Hyp .

Thus, a POCC learning rule can be translated to a quantum one with exactly the same error characteristics. The situation becomes more complicated when we generalize to truly quantum learning settings, because certain operations that are possible in the classical case are not possible in the quantum. In particular, in quantum settings, the hypothesis class consists of non-orthogonal states, which cannot be almost surely distinguished from one another. Thus, our learning rules cannot be functions of the inputs themselves, but instead can only be functions of outcomes of measurements applied to these inputs.

B. Failure of uniform convergence and ERM for PAC learnable probabilistically observed hypothesis classes

We next develop our first main results. The empirical risk minimization (ERM) rule is a cornerstone of statistical learning theory in the setting of deterministic concept classes. For a dataset $S = \{(X_i, Y_i)\}_{i=1}^m$, the empirical risk of a hypothesis $h \in \text{Hyp}$ is given by

$$\hat{R}(h, S) = \frac{1}{m} \sum_{j=1}^m \ell(h[X_j], Y_j). \quad (1)$$

In the deterministic case, a hypothesis class Hyp being PAC learnable is logically equivalent to ERM being a PAC learning rule, which is logically equivalent to it satisfying the following uniform convergence property: for any hypothesis $h \in \text{Hyp}$ and any data-generating distribution \mathcal{D} , $\mathbb{P}_{S \sim \mathcal{D}}[|R(h) - \hat{R}(h, S)| \geq \epsilon] \leq \delta$.

The ERM rule has been proposed for use as a subroutine in the quantum setting in prior work [1] and also adopted in the more recent work [4]. Both of these works give sample complexity upper bounds for this ERM rule. Our first main result is that uniform convergence and the ERM rule can fail for a POCC class Hyp , despite Hyp being PAC learnable. This

is in stark contrast to the deterministic case. We show, in our Theorem II.6, that this has further implications for the quantum setting, and thus for the tightness of the bounds in [4].

Theorem II.4 (Failure of uniform convergence and ERM for POCC classes). *There exists a POCC class Hyp that is PAC learnable but for which the ERM rule is not PAC and does not satisfy the uniform convergence property.*

Furthermore, there exists a POCC class $\widehat{\text{Hyp}}$ and a choice of \mathcal{X}, \mathcal{Y} , and \mathcal{D} for which the uniform convergence property is not satisfied, but the ERM rule is PAC.

C. Failure of uniform convergence for most finite-dimensional POVM classes

We next show that the situation regarding ERM is even worse in the quantum case. In particular, a consequence of what we show next is that the sample complexity upper bounds in [4] are infinite (i.e., vacuous) for a very large class of POVM classes that are learnable. To do so, we recall the definition of a deterministic POVM.

Definition II.5 (Deterministic POVM). A POVM $\Pi = \{\Pi_0, \Pi_1\}$ is deterministic if either $\Pi_0 = 0$ or $\Pi_1 = 0$.

That is, the outcome of a deterministic POVM is the same when used to measure any state. We note that if a POVM is not deterministic, then its outcome is statistically dependent on the state that it is used to measure.

We also define the L_1 operator norm for operators on a Hilbert space \mathcal{H} . For an operator $\Gamma : \mathcal{H} \rightarrow \mathcal{H}$, the L_1 operator norm is given by $\|\Gamma\|_{op, L_1} = \sup_{x \in \mathcal{H}} \frac{\|\Gamma x\|_1}{\|x\|_1}$. Any norm generates a topology, which allows us to talk about open and closed sets. This is used in the next theorem.

Theorem II.6 (Failure of uniform convergence of ERM for most finite-dimensional POVM classes). *Let \mathcal{X} be a subset of a finite-dimensional Hilbert space \mathcal{H} . Consider an L_1 operator norm-closed POVM hypothesis class Hyp satisfying the following conditions:*

- 1) Hyp is jointly measurable⁵.
- 2) Hyp has infinite cardinality.

Then exactly one of the following conclusions holds:

- 1) *Uniform convergence for ERM does not hold for Hyp , and ERM is not PAC.*
- 2) *The only points of accumulation of Hyp are deterministic POVMs.*

Theorem II.6 effectively says that an infinite-cardinality (but possibly finite-dimensional) POVM class can only enjoy the uniform convergence property for ERM if it ‘‘clusters’’ around deterministic measurements. The only deterministic measurements are the ones whose outcomes do not depend on the states being measured. This implies that ERM is not a useful learning rule for a rich enough set of POVM classes. Since ERM was a core subroutine of [4], this provides useful

⁵A set of POVMs is called jointly measurable if there exists a *root* POVM Π such that every $h \in \text{Hyp}$ can be simulated on a state ρ by measuring ρ with Π and then classically post-processing the outcome.

insight on prior work: in particular, in that work, a sample complexity upper bound for the ERM rule is given in the case where one can find a finite-cardinality jointly measurable partition. Our theorem above implies that this upper bound must be ∞ unless almost all of the hypotheses are close to deterministic (and, thus, independent of the input state). In our subsequent theorem, we will show that the upper bound of infinity is, in finite-dimensional cases, loose.

III. MAIN RESULTS: EVERY FINITE-DIMENSIONAL POVM CLASS IS LEARNABLE

In this section, we give a complete characterization of learnability of POVM classes in the case where \mathcal{X} is a (**possibly infinite-cardinality**) subset of the set of density operators on a finite-dimensional Hilbert space \mathcal{H} . We call a POVM class defined on \mathcal{X} a finite-dimensional POVM class. It turns out that *every finite-dimensional POVM class is learnable* – Theorem III.3. To state the sample complexity bound, we introduce the total variation covering number of a POVM hypothesis class.

Definition III.1 (Total variation distance between POVMs). Let Π_1, Π_2 be two POVMs with common domain \mathcal{X} . We define the total variation distance between Π_1, Π_2 as follows:

$$d_{TV}(\Pi_1, \Pi_2) = \sup_{x \in \mathcal{X}} d_{TV}(\text{Out}(\Pi_1, x), \text{Out}(\Pi_2, x)), \quad (2)$$

where $\text{Out}(\Pi, x)$ denotes the random outcome of the POVM Π on the mixed state x .

Definition III.2 (Total variation covering number of Hyp). Let Hyp be a POVM hypothesis class. We define an ϵ - d_{TV} covering of Hyp to be a subset $\text{Hyp}' \subseteq \text{Hyp}$ such that for any $h \in \text{Hyp}$, there exists $h' \in \text{Hyp}'$ such that $d_{TV}(h, h') \leq \epsilon$. We then define the ϵ -total variation covering number of Hyp to be the infimum cardinality over all possible ϵ - d_{TV} coverings of Hyp.

Theorem III.3 (Every finite-dimensional POVM class is learnable). *Let the span of the domain \mathcal{X} be a finite-dimensional subspace of the space of density operators on a Hilbert space \mathcal{H} . Let Hyp be a POVM class all of whose POVMs are defined on \mathcal{X} . Then Hyp is PAC learnable with the following sample complexity:*

$$n_{\text{Hyp}}(\epsilon, \delta) \leq \sum_{r=1}^N \frac{8}{\epsilon^2} \log \frac{2N}{\delta} = \frac{8N}{\epsilon^2} \log \frac{2N}{\delta}, \quad (3)$$

where $N < \infty$ is the $\frac{\epsilon}{4}$ -TV covering number of Hyp.

In the worst case, the covering number in Theorem III.3 can be exponential in the dimension of the Hilbert space. However, hypothesis classes of interest, where the POVMs have constrained structure, have a much smaller covering number. Additionally, in certain cases, one can take advantage of joint measurability in order to tighten this bound.

The above theorem provides infinitely many examples of POVM classes that are learnable. Furthermore, this class of examples includes ones such that the sample complexity

upper bounds given in [4] were infinite. Therefore, this is a substantial improvement on the previous results.

Interestingly, the proof involves concocting a learning rule that uses ERM, but in a different manner from prior work – in particular, on a *smoothing* of the original hypothesis class. This approach only works for the finite-dimensional case, necessitating yet another learning rule for our subsequent results – included in the journal version of this paper.

The proof of Theorem III.3 consists of the following steps:

- 1) We show that the total variation distance has the property that any two hypotheses within distance γ of each other have expected risks within γ of each other. Throughout, we choose $\gamma = \epsilon/4$.
- 2) We show that for every γ , the d_{TV} γ -covering number of a finite-dimensional POVM class is finite: i.e., it can be covered by finitely many d_{TV} balls of radius $< \gamma$.
- 3) Using the finiteness of covering numbers, we define the smoothing of the hypothesis class by a given γ -covering, which is a hypothesis class consisting of the centers of the balls in the covering. This class is necessarily finite-cardinality.
- 4) By previous results in the literature [1], the smoothed class is agnostically PAC learnable via ERM because it is finite-cardinality. The output of a $(\epsilon/4, \delta)$ -PAC learning rule on this hypothesis class has true risk within $\epsilon/2$ of the minimum possible within the smoothed class. This minimum has, by our result on the total variation metric, a true risk that is within $\epsilon/2$ of the infimum of possible true risks in the original hypothesis class. Thus, the hypothesis returned by the learning rule on the smoothed class has true risk within ϵ of the infimum for the original class, with probability at least $1 - \delta$.

IV. CONCLUSION

We have studied learnability of POVM hypothesis classes and shown that the standard ERM rule fails to satisfy the uniform convergence property and to be a universal learning rule. This is in contrast to the classical setting of deterministic hypotheses, and it illustrates limitations of certain prior results. Nonetheless, we showed that all finite-dimensional POVM classes are learnable via a modification of ERM, and we provided quantitative sample complexity bounds for this learning rule. There are various possible extensions of our work: for instance, a characterization of the sample complexity of learning a hypothesis class in terms of its Hilbert space geometry would be of interest. Additionally, our learning rule only makes *separable* measurements. In quantum hypothesis testing, where the goal is to distinguish between two *known* states with minimal error probability from m copies of one of them, block measurements have a provable advantage in terms of sample complexity. It would be interesting to understand whether this phenomenon holds in the learning setting.

ACKNOWLEDGMENT

This research is funded by NSF CCF grant number #2212327.

REFERENCES

- [1] M. Heidari, A. Padakandla, and W. Szpankowski, "A theoretical framework for learning from quantum data," in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 1469–1474.
- [2] M. Heidari, A. Grama, and W. Szpankowski, "Toward physically realizable quantum neural networks," 2022. [Online]. Available: <https://arxiv.org/abs/2203.12092>
- [3] Y. Wu, B. Wu, J. Wang, and X. Yuan, "Quantum phase recognition via quantum kernel methods," *Quantum*, vol. 7, p. 981, apr 2023. [Online]. Available: <https://doi.org/10.22331/q-2023-04-17-981>
- [4] 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. 11 305–11 319. [Online]. Available: <https://proceedings.mlr.press/v151/padakandla22a.html>
- [5] S. Arunachalam and R. De Wolf, "Optimal quantum sample complexity of learning algorithms," *J. Mach. Learn. Res.*, vol. 19, no. 1, p. 2879–2878, Jan. 2018.
- [6] M. Wilde, *Quantum information theory*. Cambridge, UK: Cambridge University Press, 2017. [Online]. Available: <https://arxiv.org/abs/1106.1445>
- [7] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014.
- [8] S. Arunachalam and R. de Wolf, "Guest column: A survey of quantum learning theory," *SIGACT News*, vol. 48, no. 2, p. 41–67, jun 2017. [Online]. Available: <https://doi.org/10.1145/3106700.3106710>
- [9] S. Arunachalam, A. B. Grilo, and H. Yuen, "Quantum statistical query learning," *arXiv preprint arXiv:2002.08240*, 2020.
- [10] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, "Sample-efficient learning of interacting quantum systems," *Nature Physics*, pp. 1745–2481, 2021. [Online]. Available: <https://www.nature.com/articles/s41567-021-01232-0>
- [11] R. O'Donnell and J. Wright, "Efficient quantum tomography," in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 899–912. [Online]. Available: <https://doi.org/10.1145/2897518.2897544>
- [12] —, "Efficient quantum tomography ii," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, ser. STOC 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 962–974. [Online]. Available: <https://doi.org/10.1145/3055399.3055454>
- [13] 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*, ser. STOC '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 913–925. [Online]. Available: <https://doi.org/10.1145/2897518.2897585>
- [14] S. Aaronson, "The learnability of quantum states," *Proceedings of The Royal Society A Mathematical Physical and Engineering Sciences*, vol. 463, 08 2006.
- [15] J. Rehacek and M. Paris, *Quantum state estimation*, ser. Lecture notes in physics. Berlin: Springer, 2004. [Online]. Available: <https://cds.cern.ch/record/1391369>
- [16] M. H. Waseem, F. e Ilahi, and M. S. Anwar, "Quantum state tomography," in *Quantum Mechanics in the Single Photon Laboratory*, ser. 2053-2563. IOP Publishing, 2020, pp. 6–1 to 6–20. [Online]. Available: <http://dx.doi.org/10.1088/978-0-7503-3063-3ch6>
- [17] J. Altepeter, D. James, and P. Kwiat, "Qubit quantum state tomography," *Lecture Notes in Physics*, vol. 649, pp. 113–145, 2004.
- [18] S. Aaronson, "Shadow tomography of quantum states," in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, ser. STOC 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 325–338. [Online]. Available: <https://doi.org/10.1145/3188745.3188802>
- [19] S. Aaronson, X. Chen, E. Hazan, and S. Kale, "Online learning of quantum states," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, ser. NIPS'18. Red Hook, NY, USA: Curran Associates Inc., 2018, p. 8976–8986.
- [20] 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.
- [21] A. Atici and R. A. Servedio, "Improved bounds on quantum learning algorithms," *Quantum Information Processing*, vol. 4, no. 5, pp. 355–386, oct 2005. [Online]. Available: <https://doi.org/10.1007/s2Fs11128-005-0001-2>
- [22] M. Roget, G. D. Molfetta, and H. Kadri, "Quantum perceptron revisited: Computational-statistical tradeoffs," in *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022. [Online]. Available: <https://openreview.net/forum?id=BuOcoU8jqxq>
- [23] H.-C. Cheng, M.-H. Hsieh, and P.-C. Yeh, "The learnability of unknown quantum measurements," *Quantum Info. Comput.*, vol. 16, no. 7–8, p. 615–656, May 2016.
- [24] 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.
- [25] M. J. Kearns and R. E. Schapire, "Efficient distribution-free learning of probabilistic concepts," in *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, ser. SFCS '90. USA: IEEE Computer Society, 1990, p. 382–391 vol.1. [Online]. Available: <https://doi.org/10.1109/FSCS.1990.89557>