

On the Hardness of PAC-learning Stabilizer States with Noise

Aravind Gollakota and Daniel Liang

Department of Computer Science
University of Texas at Austin
January 31, 2022

We consider the problem of learning stabilizer states with noise in the Probably Approximately Correct (PAC) framework of Aaronson [1] for learning quantum states. In the noiseless setting, an algorithm for this problem was recently given by Rocchetto [41], but the noisy case was left open. Motivated by approaches to noise tolerance from classical learning theory, we introduce the Statistical Query (SQ) model for PAC-learning quantum states, and prove that algorithms in this model are indeed resilient to common forms of noise, including classification and depolarizing noise. We prove an exponential lower bound on learning stabilizer states in the SQ model. Even outside the SQ model, we prove that learning stabilizer states with noise is in general as hard as Learning Parity with Noise (LPN) using classical examples. Our results position the problem of learning stabilizer states as a natural quantum analogue of the classical problem of learning parities: easy in the noiseless setting, but seemingly intractable even with simple forms of noise.

1 Introduction

A fundamental task in quantum computing is that of learning a description of an unknown quantum state ρ . Traditionally this is formalized as the problem of quantum state tomography, where we are granted the ability to form multiple copies of ρ and take arbitrary measurements, and must learn a state σ that is close to ρ in trace distance. In an influential work, Aaronson [1] introduced the “Probably Approximately Correct” (PAC) framework from computational learning theory [43] as an alternative perspective on this problem. Here the key innovation is that instead of learning ρ in an absolute metric (such as trace distance), we only wish to learn it with respect to a pre-specified distribution on measurements. This requirement is considerably weaker than that of full tomography.

Concretely, let \mathcal{E} denote the space of two-outcome n -qubit measurements E , where E (corresponding to the POVM $\{E, I - E\}$) accepts a state ρ with probability $\text{Tr}(E\rho)$ and rejects it otherwise. Let D be a distribution on \mathcal{E} . We are given the ability to sample and

Aravind Gollakota: aravindg@cs.utexas.edu

Daniel Liang: dliang@cs.utexas.edu

perform random measurements from D , i.e. to form a training set $\{(E_i, Y_i)\}_{i \in [m]}$ where each $E_i \sim D$ and Y_i is the outcome of measuring ρ using E_i . Our goal is to learn a state σ such that with high probability,

$$\mathbb{E}_{E \sim D} \left[(\text{Tr}(E\sigma) - \text{Tr}(E\rho))^2 \right] \leq \epsilon.$$

Usually we also have some basic prior information about ρ , such as knowledge of a class \mathcal{F} such that $\rho \in \mathcal{F}$. In the terminology of learning theory, this would be called “learning the class \mathcal{F} .” It is important to stress that in this framework, while the data *arises from* a quantum state, it is entirely classical in form and representation.

When ρ is an arbitrary n -qubit state, [1] showed that $O(n)$ “training examples”, i.e. measurement-outcome pairs $(E, \text{Tr}(E\rho))$,¹ suffice to statistically determine a state σ that approximates ρ in this sense (suppressing the dependence on ϵ for simplicity). This is in contrast to full state tomography, which requires an exponential number of measurements or copies of ρ [28, 36]. Note, however, that here σ is determined purely in a statistical or information-theoretic sense, and for a long time no efficient algorithms were known for actually computing σ in settings of interest.

Recently, Rocchetto [41] gave an efficient algorithm for the case where ρ is known to be a stabilizer state, and D is any distribution on Pauli measurements. The stabilizer states are an important class of states in quantum computing due to their efficient classical simulability [3, 26] and their foundational role in quantum error-correction [25]. A major question left open by [41] is: are stabilizer states also efficiently learnable in noisy settings?

Motivated by this question, we introduce to the quantum setting a well-known tool for noise-resilient classical PAC learning, the *statistical query* (SQ) model, and define the problem of SQ-learning quantum states. In this model, rather than receiving labeled measurement-outcome examples of the form $(E, \text{Tr}(E\rho))$, the learner is only allowed to make statistical queries to an oracle, and otherwise its goal remains the same. A statistical query is described by a function $\varphi : \mathcal{E} \times \{-1, 1\} \rightarrow [-1, 1]$ and a tolerance $\tau > 0$, and the oracle responds to the query with $\mathbb{E}[\varphi(E, Y)] \pm \tau$, where the expectation is taken over the random draw of $E \sim D$ and Y , the random outcome of measuring ρ using E . One can think of this as modeling an experimental setup that is unable to report individual measurement outcomes, but is nevertheless able to estimate expectation values to any desired accuracy. Importantly, an algorithm that is able to work in this restricted setting automatically gains tolerance to several kinds of noise.

The SQ model was originally introduced by Kearns [32] in the setting of Boolean function classes, and has since grown into a highly influential model (see [23, 40] for surveys). The model is known to have the following properties:

- It is a natural restriction of the PAC model that nevertheless captures most known PAC algorithms for a wide range of common classes [30, 40].
- SQ algorithms are naturally resistant to mild forms of noise in the labels, such as “classification noise”, where the label for each training example is flipped with some constant probability [32].

¹Or more realistically (E, Y) , where Y is the random outcome of measuring ρ using E .

- It is the most realistic learning model for which strong, unconditional lower bounds are known for many basic classes. Indeed, there is a considerable literature on this topic, with lower bounds usually proven using the so-called SQ-dimension and its generalizations [15, 22, 40].
- SQ algorithms are naturally implementable in a way that satisfies differential privacy of the training data, and indeed are the main examples of realistic differentially private learning algorithms [17, 21].

Given all of these properties, it is natural to wonder whether the SQ model has something to bring to quantum learnability, with a particular eye towards noise tolerance. In this work we show (among other results) that for stabilizer states this approach cannot work: SQ-learning stabilizer states is exponentially hard, and in general, learning stabilizer states with noise is as hard as the well-known Learning Parity with Noise (LPN) problem.

Theorem 1.1. *Let D denote the uniform distribution on Pauli measurements. Any SQ algorithm for learning n -qubit stabilizer states under D up to error $2^{-O(n)}$ (i.e. to significantly outperform the maximally mixed state) requires $2^{\Omega(n^2)}$ queries even when tolerance is $2^{-O(n^2)}$.*

Define a *parity measurement* to be a Pauli measurement of the form $E_x = \frac{P_x + I}{2}$ for some $x \in \{0, 1\}^n$, where $P_x = \sum_{y \in \{0, 1\}^n} \chi_x(y) |y\rangle\langle y|$ and $\chi_x(y) = (-1)^{x \cdot y \bmod 2}$. They are so named since for any computational basis state $|y\rangle\langle y|$, $\text{Tr}(E_x |y\rangle\langle y|) = x \cdot y \bmod 2$. The following theorems hinge on the observation (stated as Proposition 4.10) that parities can be very naturally embedded within the problem of learning stabilizer states under distributions on parity measurements.

Theorem 1.2. *Let D' denote the uniform distribution on parity measurements. Any SQ algorithm for learning n -qubit stabilizer states under D' even up to constant error (say $1/3$) requires $2^{\Omega(n)}$ queries even when tolerance is $2^{-O(n)}$.*

Theorem 1.3. *Let D' be as above. Learning n -qubit stabilizer states under D' with classification noise at rate η is at least as hard as the classical problem of Learning Parity with Noise (LPN) at rate η .*

(These theorems are formally stated as Corollaries 4.7, 4.11 and 4.12 respectively.)

Our results position the problem of learning stabilizer states as a quantum analogue of the important classical problem of learning parities.² In both cases there are simple “algebraic” learning algorithms for the noiseless setting, and the problem seems to become intractable with even the simplest kinds of noise. The algorithm of Rocchetto [41] thus joins a small class of PAC algorithms that do not fall into the SQ model, and hence do not admit any straightforward algorithms in noisy settings. In our view, this frames learning stabilizer states with noise as one of the more compelling problems on the frontier of learning quantum states with noise.

Another interpretation of our results is that they highlight limitations of the PAC framework of Aaronson [1]: insofar as this framework reduces the problem of learning quantum states to an essentially classical problem, it also inherits longstanding problems from classical learning theory.

²This recalls the way in which stabilizer codes are a quantum analogue of classical parity check codes.

We also hope that our introduction of the SQ model to quantum state learning will be of independent interest and help spur new ideas in this area.

We now detail the rest of our contributions and lay out the organization of this paper:

- In Section 2, we formally define the problem of SQ-learning quantum states and extend the notion of the SQ-dimension to this setting, building on recent work that formally analyzed the SQ-dimension as applicable to the p-concept setting [24].
- In Section 3, we show that SQ algorithms for learning quantum states are indeed resistant to mild forms of noise, including classical classification noise as well as quantum channels with bounded noise (such as depolarizing noise).
- In Section 4, we give exponential SQ lower bounds on learning stabilizer states. Under the uniform distribution on Pauli measurements, we show (Corollary 4.7) that it requires exponentially many queries in order to improve on the maximally mixed state's performance. Under a different natural distribution on Pauli measurements, namely the uniform distribution over parity measurements, we show (Corollaries 4.11, 4.12) that learning stabilizer states with noise is as hard as learning parities with noise.
- In Section 5, by way of positive results, we give SQ algorithms for the simple setting of learning product states. We describe an SQ algorithm for learning product states under Haar-random single-qubit measurements, and show that it allows one to perform tomography on the individual qubits.
- In Section 6, we relate SQ learning to a form of differential privacy for quantum state learners. This form of differential privacy has recently been studied by [11].

1.1 Related work

The problem of learning quantum states via state tomography has a long history in quantum computing, culminating in the celebrated optimal algorithms of O'Donnell and Wright [36, 37] and Haah et al. [28]. We operate in the alternative PAC framework introduced by Aaronson [1]. In recent years, this framework has been extended to the online setting [6] as well as verified in experimental setups [42]. To our knowledge, the only known computationally efficient PAC learners for supervised learning of a commonly-considered class of states are the algorithm of Rocchetto [41] for learning stabilizer states, as well as that of Yoganathan [45] for other classes of states whose generating circuits can be efficiently classically simulated and inverted, including low Schmidt rank states. While the focus of this paper is on stabilizer states, we remark that Yoganathan's algorithm for low Schmidt rank states also involves solving a system of polynomial equations in the examples, and hence would also not admit any straightforward SQ implementation. Cheng et al. [19] frame the problem of PAC-learning unknown quantum measurements under a distribution of states as a dual problem to PAC-learning an unknown state, and are able to recover Aaronson's main sample complexity bound using a classical proof.

Other examples of efficient learning algorithms not lying strictly within the PAC formalism include Aaronson and Grewal [4] in the unsupervised setting for non-interacting fermions and Montanaro [35] for improved tomography of stabilizer states. Lai and Cheng [34] extend

Montanaro’s ideas for stabilizer states and use it to learn outputs of Clifford circuits with a single layer of T gates.

Recent work by Arunachalam et al. [11] extends work by Bun et al. [18] to the quantum setting, and relates differentially private (DP) learning of quantum states to one-way communication, online learning, and other models. We show in Section 6 that our notion of SQ learnability implies their notion of DP learnability, and hence by their results also implies finite sequential fat-shattering dimension, online learnability, and “quantum stability.”

We note that the problem of PAC-learning quantum states is very different from the problem of PAC-learning Boolean functions using quantum representations of data, as considered in a recent active line of work [8, 9]. In particular, the model of SQ-learning that we introduce is unrelated to a recent notion of SQ-learning of *Boolean* functions using quantum representations [10]. When one is given quantum samples of Boolean or integer-valued functions, there have been important results on learning in the presence of noise, showing that both Learning Parity with Noise (LPN) [20] and Learning With Errors (LWE) [27] are tractable in this setting.

2 Preliminaries

Notation and terminology. We use ρ to refer to the density matrix of an n -qubit quantum mixed state, representable as a $2^n \times 2^n$ PSD operator of trace 1. (The number of qubits will be n throughout and suppressed from the notation.) A pure state is a quantum state with rank 1. Let \mathcal{E} denote the space of two-outcome n -qubit measurements E (corresponding to the POVM $\{E, I - E\}$), which accept a state ρ with probability $\text{Tr}(E\rho)$. We will view the measurement outcomes themselves as $\{-1, 1\}$ -valued, so that the outcome of measuring ρ using E is a random variable Y that is 1 with probability $\text{Tr}(E\rho)$ and -1 otherwise. Define $f_\rho : \mathcal{E} \rightarrow [-1, 1]$ to be the conditional mean function

$$f_\rho(E) = \mathbb{E}[Y|E] = 1 \cdot \text{Tr}(E\rho) + (-1) \cdot (1 - \text{Tr}(E\rho)) = 2 \text{Tr}(E\rho) - 1.$$

We will often identify a state ρ with its behavior with respect to two-outcome measurements, namely with the function f_ρ , and use the notation $Y \sim f_\rho(E)$ to mean that $Y \in \{-1, 1\}$ is the random measurement outcome satisfying $\mathbb{E}[Y|E] = f_\rho(E)$. In learning theoretic terms, this means f_ρ describes a *probabilistic concept*, or p-concept, on the space \mathcal{E} . A p-concept on a domain \mathcal{X} is a classification rule that assigns random $\{-1, 1\}$ -valued labels to each point in \mathcal{X} according to a fixed conditional mean function; we always identify the p-concept with its conditional mean function. Given a set \mathcal{F} of quantum states, we use \mathcal{F} to also mean the class of associated p-concepts, with the meaning clear from context.

Given a distribution D over \mathcal{E} , we will often regard functions $f_\rho, f_\sigma : \mathcal{E} \rightarrow \mathbb{R}$ as members of the L^2 space $L^2(D, \mathcal{E})$, with the inner product given by $\langle f_\rho, f_\sigma \rangle_D = \mathbb{E}_{E \sim D}[f_\rho(E)f_\sigma(E)]$, and the norm given by $\|f_\rho\|_D = \sqrt{\langle f_\rho, f_\rho \rangle_D} = \sqrt{\mathbb{E}_{E \sim D}[(f_\rho(E))^2]}$.

We use $[n]$ to refer to the set of indices $\{1, \dots, n\}$. Given a set $S \subseteq [n]$, we will use χ_S to refer to the parity on S , defined as a function from $\{-1, 1\}^n \rightarrow \{-1, 1\}$ by $\chi_S(x) = \prod_{i \in S} x_i$. Given $x, y \in \{0, 1\}^n$, we will sometimes also use $\chi_x(y) = (-1)^{x \cdot y \bmod 2}$.

2.1 Learning models

We begin by formally defining the problem of PAC-learning a quantum state.

Definition 2.1 (PAC-learnability of quantum states, [1]). Let \mathcal{F} be a class of n -qubit quantum states. Let D be a distribution over \mathcal{E} . We say \mathcal{F} is PAC-learnable up to squared loss ϵ with respect to D if there exists a learner that, given sample access to labeled examples (E, Y) for $E \sim D, Y \sim f_\rho(E)$ for an unknown $\rho \in \mathcal{F}$, is able to output a state σ satisfying

$$\mathbb{E}_{E \sim D} \left[(f_\sigma(E) - f_\rho(E))^2 \right] \leq \epsilon.$$

The number of examples used by the learner is called its sample complexity.

We note that this is a slight modification of the original definition in [1], stated directly in terms of squared loss since this is the view that will be convenient for us. PAC learners are also allowed to fail with some probability δ , but for simplicity we will ignore this in this paper. A learner that only succeeds with some constant probability can easily be amplified to succeed with probability $1 - \delta$ using standard confidence amplification procedures.

With PAC learners, one may speak of both computational efficiency (overall running time) and statistical or information-theoretic efficiency (sample complexity). The original result of Aaronson [1] described a computationally inefficient algorithm for learning arbitrary states that nevertheless had $O(n)$ sample complexity. An *efficient* PAC learner is one that is computationally efficient, i.e. runs in polynomial time, and hence also draws at most polynomially many examples (each draw is considered as taking one unit of time).

We now introduce the following natural extension of these definitions to the SQ setting. In both cases, we operate in the so-called distribution-specific setting, where the learner is assumed to have knowledge of the distribution D .

Definition 2.2 (SQ-learnability of quantum states). Let \mathcal{F} be a class of n -qubit quantum states. Let D be a distribution over \mathcal{E} . An SQ oracle for an unknown state $\rho \in \mathcal{F}$ is an oracle that accepts a query and a tolerance, (φ, τ) , where $\varphi : \mathcal{E} \times \{-1, 1\} \rightarrow [-1, 1]$ and $\tau > 0$, and responds with y such that

$$\left| y - \mathbb{E}_{E \sim D, Y \sim f_\rho(E)} [\varphi(E, Y)] \right| \leq \tau.$$

We say \mathcal{F} is SQ-learnable up to squared loss ϵ if there is a learner that, given only queries to the SQ oracle for an unknown $\rho \in \mathcal{F}$, is able to output a state σ satisfying

$$\mathbb{E}_{E \sim D} \left[(f_\sigma(E) - f_\rho(E))^2 \right] \leq \epsilon.$$

The number of queries used by the learner is called its query complexity.

An SQ learner is considered *efficient* if it uses polynomially many queries and its queries all have tolerance $\tau \geq 1/\text{poly}(n)$.

2.2 SQ lower bounds for p-concepts

One of the chief features of the classical SQ model is the possibility of proving unconditional lower bounds on learning a class \mathcal{C} in terms of its so-called statistical dimension. The quantum setting that we work in, where we identify a state ρ with the p-concept f_ρ , becomes a special case of the SQ model for learning p-concepts. Building on recent work [24] that formally proved SQ lower bounds for p-concepts, we extend this framework to the quantum setting. Let \mathcal{X} denote an arbitrary domain (for us, \mathcal{X} will be \mathcal{E} , while in the classical setting, \mathcal{X} is usually \mathbb{R}^n).

Definition 2.3 (Statistical dimension). Let D be a distribution on \mathcal{X} , and let \mathcal{C} be a class of functions from \mathcal{X} to \mathbb{R} . The *average (un-normalized) correlation* of \mathcal{C} is defined to be $\rho_D(\mathcal{C}) = \frac{1}{|\mathcal{C}|^2} \sum_{c, c' \in \mathcal{C}} |\langle c, c' \rangle_D|$. The *statistical dimension on average* at threshold γ , $\text{SDA}_D(\mathcal{C}, \gamma)$, is the largest d such that for all $\mathcal{C}' \subseteq \mathcal{C}$ with $|\mathcal{C}'| \geq |\mathcal{C}|/d$, $\rho_D(\mathcal{C}') \leq \gamma$.

Theorem 2.4 ([24], Cor. 4.6). *Let D be a distribution on \mathcal{X} , and let \mathcal{C} be a p -concept class on \mathcal{X} . Say our queries are of tolerance τ , the final desired squared loss is ϵ , and that the functions in \mathcal{C} satisfy $\|c\|_D \geq \beta$ for all $c \in \mathcal{C}$. For technical reasons, we require $\tau \leq \epsilon$, $\epsilon^2 \leq \beta/3$. Then learning \mathcal{C} up to squared loss ϵ (we may pick ϵ as large as $\sqrt{\beta/3}$) requires at least $\text{SDA}_D(\mathcal{C}, \tau^2)$ queries of tolerance τ .*

We remark that the way to interpret such a lower bound is as follows: if the SQ learner's queries have tolerance at least τ , then at least $\text{SDA}_D(\mathcal{C}, \tau^2)$ queries are required. That is, one must *either* use small tolerance *or* many queries.

The following lemma will be convenient in order to bound the SDA when we have bounds on pairwise correlations.

Lemma 2.5 ([24], Lemma 2.6). *Let D be a distribution on \mathcal{X} , and let \mathcal{C} be a p -concept class on \mathcal{X} such that for all $c, c' \in \mathcal{C}$ with $c \neq c'$, $|\langle c, c' \rangle_D| \leq \gamma$, and for all $c \in \mathcal{C}$, $\|c\|_D^2 \leq \kappa$. Then for any $\gamma' > 0$, $\text{SDA}(\mathcal{C}, \gamma + \gamma') \geq |\mathcal{C}| \frac{\gamma'}{\kappa - \gamma}$.*

2.3 The problem of learning parities

One of the most basic problems in classical learning theory is that of learning the concept class of parity functions. Let the domain be $\mathcal{X} = \{-1, 1\}^n$, and for any subset $S \subseteq [n]$, define $\chi_S(x) = \bigoplus_{i \in S} x_i = \prod_{i \in S} x_i$ to be the parity on S . Here $x_i \oplus x_j = x_i x_j$ is simply the XOR when bits are represented by $\{-1, 1\}$. Let D be any distribution on \mathcal{X} . We say a learner is able to learn parities under D if given access to labeled examples $(x, \chi_S(x))$ where $x \sim D$ and $S \subseteq [n]$ is unknown (or, in the SQ setting, given access to the corresponding SQ oracle), and for any error parameter ϵ , it is able to output a function h such that $\mathbb{P}_{x \sim D}[h(x) \neq \chi_S(x)] \leq \epsilon$.

The problem of learning parities displays a striking phase transition in going from the noiseless to the noisy setting. Given noiseless labeled examples, the problem of recovering the right parity is simply a question of solving linear equations over \mathbb{F}_2 , and can be done using Gaussian elimination by a PAC learner using only $\Theta(n)$ examples. With just a little noise, however, the problem seems to become intractable. Perhaps the simplest noise model one can consider is the classification noise model, where every example has its label flipped with some constant probability η (known as the noise rate). Learning parities under classification noise is the basis of the famous *Learning Parity with Noise (LPN)* problem. Formally, the search version of LPN with noise rate η is precisely the problem of learning parities under the uniform distribution on \mathcal{X} and with classification noise at rate η . Usually one also has the additional knowledge that the true target χ_S (the ‘‘secret’’) is picked uniformly at random from the set of all parities. This problem is widely conjectured to be hard, including for quantum algorithms, and is even used as a basis for cryptography (see [38] for a survey). The best-known algorithm in the PAC setting runs in slightly subexponential time [16].

Since SQ learners are naturally tolerant of classification noise, one would expect that there are no SQ learners for parities under the uniform distribution, and indeed, this is one of the foundational results in the SQ literature.

Theorem 2.6 ([32]). *Any SQ learner requires $2^{\Omega(n)}$ queries (even using tolerance $2^{-O(n)}$) to learn parities under the uniform distribution on $\{-1, 1\}^n$ even up to constant error (say $1/3$).*

Thus we see that simple Gaussian elimination is an example of an efficient PAC learner that is not SQ. This establishes a characteristic limitation of SQ algorithms: while they include a wide range of common algorithms, they do not include algorithms that depend entirely on “algebraic” structure.

It is worth emphasizing that this discussion has considered learning parities with a *classical* representation of the data. When given a *quantum* representation of the data, as in the quantum “example state” $|\psi\rangle = 2^{-n/2} \sum_{x \in \{0,1\}^n} |x, \chi_S(x)\rangle$ (taking the distribution over the domain to be uniform), the task becomes easy even with noise [20]. This is because we can now use Hadamard gates to implement a Boolean Fourier transform à la the famous Bernstein–Vazirani algorithm [14].

3 Noise-tolerant SQ learning

One of the prime features of classical SQ learning is its inherent noise tolerance. From an intuitive standpoint, certain common stochastic noise models are systematic enough that their effects *in expectation* can be predicted in advance, and hence either be corrected for or bounded. Slightly more precisely, the query expectations of a noisy state are often related in simple ways to the query expectations on a noiseless state, so that the latter can be recovered from the former. We mainly consider three such noise models here: (a) classical classification noise and malicious noise, (b) quantum depolarizing noise, and (c) more general quantum channels with bounded noise.

3.1 Classification and malicious noise

Classification noise [7] and malicious noise [33, 44] are two classical Boolean noise models that SQ algorithms are able to handle. In the classification noise model, every example’s label is flipped with probability η (known as the noise rate). The malicious noise model is a stronger form of noise where for any given example, with probability $1 - \eta$, the label is reported correctly, but with probability η both the point and its label may be arbitrary (and adversarially selected based on the learner’s behavior so far). We note that these models are well-defined even in the p-concept setting and hence for quantum states, and simply introduce further randomness into the label. The following results were originally stated for Boolean functions but readily extend to p-concepts.

Theorem 3.1 ([32]). *Let \mathcal{C} be a p-concept class learnable under distribution D in the SQ model up to error ϵ using q queries of tolerance τ . Then for any constant $0 < \eta < 1/2$, even with respect to an SQ oracle with classification noise at rate η (i.e., one that computes expectations with classification noise), \mathcal{C} is learnable up to ϵ using $O(q)$ queries of tolerance $O(\tau(1 - 2\eta))$. If the learner is given noisy training examples as opposed to access to a noisy SQ oracle, then $\tilde{O}(\frac{q}{\text{poly}(\tau(1-2\eta))})$ noisy examples suffice.*

Theorem 3.2 ([12]). *Let \mathcal{C} be a p-concept class learnable under distribution D in the SQ model up to error ϵ using q queries of tolerance τ . An SQ oracle with malicious noise at rate η is one that computes query expectations with respect to a distribution $(1 - \eta)f(D) + \eta Q$, where*

$f(D)$ denotes the true labeled distribution (x, y) for $x \sim D, y \sim f(x)$ (f being the unknown target p -concept), and Q is an arbitrary and adversarially selected distribution on $\mathcal{X} \times \{-1, 1\}$. If $\eta = \tilde{O}(\epsilon)$ and $\eta < \tau$, then even with respect to an SQ oracle with malicious noise at rate η , \mathcal{C} is learnable up to ϵ using $O(q)$ queries of tolerance $\tau - \eta$. If the learner is given noisy training examples as opposed to access to a noisy SQ oracle, then \mathcal{C} is learnable (with constant probability) using $\tilde{O}(\frac{q}{\text{poly}(\tau - \eta)})$ noisy examples suffice. (More efficient implementations are also available in some special cases).

The proofs of both theorems are similar: one first relates the noisy query expectations to the true expectations, and then argues that when using a suitably small tolerance (or sufficiently many examples) the effects of the noise can be corrected for (within information theoretic limits).

3.2 Depolarizing noise

Depolarizing noise acts on quantum states by shifting them closer to the maximally mixed state. One can consider a setting where it acts on an entire n -qubit state at once, as well as one where it acts independently on each individual qubit. We will consider the former.

Definition 3.3 (Depolarizing noise). Let ρ be an arbitrary n -qubit state. Then depolarizing noise at rate η ($0 < \eta < 1$) acts on this state by transforming it into $\Lambda_\eta(\rho) = (1 - \eta)\rho + \eta(I/2^n)$.

Theorem 3.4. Let $0 < \eta < 1$ be any constant, and let Λ_η denote the depolarizing channel at noise rate η . Let \mathcal{C} be a class of n -qubit quantum states and D be a distribution on \mathcal{E} , the space of two-outcome measurements on such states. Let L be an SQ learner capable of learning \mathcal{C} under D using q queries of tolerance τ . There there exists a learner L' such that for any $\rho \in \mathcal{C}$, L' is capable of learning ρ under D using q queries of tolerance $\tau(1 - \eta)$ given only SQ access to $\Lambda_\eta(\rho)$ as well as sampling access to D .

Proof. For simplicity, we will assume that we know the noise rate η exactly. (So long as we have an upper bound on η , then by a standard “grid search” argument due to [32], we can estimate η sufficiently closely simply by trying out many different values. Briefly: if say we try out $\eta = 0, \delta, 2\delta, \dots, 1$ ($1/\delta$ values in all), then one of these will be within $\delta/2$ of the true η . The algorithm when run with this guess for η will produce a good hypothesis. By taking $\delta = O(\tau(1 - \eta)^2)$ and testing all $1/\delta$ hypotheses produced by our guesses for η on a sufficiently large validation set, we can ensure the best one will perform and generalize well.)

Let $\rho \in \mathcal{C}$ be the unknown target. Observe that for any $E \in \mathcal{E}$, by linearity,

$$f_{\Lambda_\eta(\rho)}(E) = 2 \text{Tr}(E \cdot \Lambda_\eta(\rho)) - 1 = (1 - \eta)f_\rho(E) + \eta f_{I/2^n}(E).$$

Let $\varphi : \mathcal{E} \times \{-1, 1\} \rightarrow [-1, 1]$ be any query made by L . Let $\varphi[\rho]$ denote the query expectation of φ under ρ , given by $\mathbb{E}_{x \sim D} \mathbb{E}_{y \sim f_\rho(x)}[\varphi(x, y)]$. Similarly let the noisy analogue be $\varphi[\Lambda_\eta(\rho)]$. Again just by linearity,

$$\varphi[\Lambda_\eta(\rho)] = (1 - \eta)\varphi[\rho] + \eta\varphi[I/2^n].$$

The latter quantity is independent of ρ and can be estimated to arbitrary accuracy by sampling from D , allowing us to estimate $\varphi[\rho]$ as $\frac{\varphi[\Lambda_\eta(\rho)] - \eta\varphi[I/2^n]}{1 - \eta}$. So long as η is bounded away from 1, we can use a query of tolerance $\tau(1 - \eta)$ to estimate $\varphi[D_\eta(\rho)]$ (as well as $1/\text{poly}(\tau(1 - \eta))$ unlabeled examples from D to compute $\varphi[I/2^n]$), and thereby estimate $\varphi[\rho]$ to within τ . Thus we can simulate L even with depolarizing noise. \square

It is worth stressing that we are able to handle any constant noise rate $\eta \in (0, 1)$, and the price we pay is requiring the tolerance to scale as $\tau(1 - \eta)$.

3.3 Quantum channels with bounded noise

We can also consider more general kinds of quantum channels with bounded noise. As long as the queries are bounded, small amounts of noise cannot alter query expectations too much, and so can be “absorbed” into the tolerance. This is similar to classical malicious noise: since classical malicious noise at rate η only can only change query expectations by η (recall that the queries are bounded by 1), a noisy query of tolerance $\tau - \eta$ is able to simulate a noiseless query of tolerance τ . Unlike with depolarizing noise, this means we cannot handle arbitrary η ; this is an artifact of the fact that more general kinds of noise do not permit the kind of systematic correction we were able to perform for depolarizing noise.

For concreteness here we consider a noisy quantum channel Λ such that $\|\Lambda - 1_n\|_\diamond \leq \eta$, where 1_n is the identity map on n -qubit states and the norm is the diamond norm. We do not define this norm here, but its chief property for our purposes is that for any n -qubit state ρ and 2-outcome measurement E , $|\text{Tr}(E(\rho - \Lambda(\rho)))| \leq \eta$. Similar theorems can be proven with respect to other distance measures such as fidelity.

Theorem 3.5. *Let Λ be a quantum channel such that $\|\Lambda - 1_n\|_\diamond \leq \eta$, as above. Let \mathcal{C} be a class of n -qubit quantum states learnable under distribution D using q queries of tolerance $\tau > 2\eta$. Then \mathcal{C} is still learnable under noise Λ (i.e. when our queries are answered not with respect to ρ but $\Lambda(\rho)$) using q noisy queries of tolerance $\tau - 2\eta$.*

Proof. As noted, for any state ρ and measurement E , $|\text{Tr}(E(\rho - \Lambda(\rho)))| \leq \eta$. Consider any query $\varphi : \mathcal{E} \times \{-1, 1\} \rightarrow [-1, 1]$. If $\varphi[\rho]$ denotes the query expectation on a noiseless state and $\varphi[\Lambda(\rho)]$ denotes the noisy one, then a straightforward calculation shows that

$$\begin{aligned} |\varphi[\rho] - \varphi[\Lambda(\rho)]| &= \left| \mathbb{E}_{E \sim D} [(\varphi(E, 1) - \varphi(E, -1)) \text{Tr}(E(\rho - \Lambda(\rho)))] \right| \\ &\leq 2 \mathbb{E}_{E \sim D} |\text{Tr}(E(\rho - \Lambda(\rho)))| \\ &\leq 2\eta. \end{aligned}$$

Thus just by the triangle inequality, if we calculate $\varphi[\Lambda(\rho)]$ within tolerance $\tau - 2\eta$, then we also get $\varphi[\rho]$ within τ . □

3.4 General noise for distribution-free learning

So far, we’ve only considered distribution-specific learning, where the learner is only required to succeed with respect to a pre-specified distribution D . In the distribution-free case, where the learner is required to succeed no matter what D is, we now give a simple proof that any SQ algorithm for a concept class can also handle any kind of quantum noise on the state, as long as the noise is known. This is unsurprising, and at a high level, the approach simply boils down to off-loading the noise from the state to the measurement. Learning a noisy set of measurements is thus handled by distribution-free learning algorithm.

Given a quantum operation Λ , its adjoint Λ^\dagger is such that $\forall \rho, \text{Tr}(E \cdot \Lambda(\rho)) = \text{Tr}(\Lambda^\dagger(E) \cdot \rho)$ and always exists (see [39] for details on how to prove this folklore result). Let \mathcal{D} be the

distribution we are trying to learn concept class \mathcal{C} using statistical queries and let Λ be the noise applied to the quantum state. We can then define \mathcal{D}^\dagger to be the distribution $\Lambda^\dagger(E)$ where E is drawn from \mathcal{D} and by definition the traces (and thus the statistical queries) are the same when applied to ρ and $\Lambda(\rho)$ respectively. Also by definition, a distribution-free learner for \mathcal{C} would also be able to learn with distribution \mathcal{D}^\dagger .

4 Lower bounds on learning stabilizer states with noise

4.1 Stabilizer formalism

The stabilizer states are a popular class of quantum states that are used throughout quantum information for areas like quantum error correction, classical simulation of quantum mechanics, and quantum communication. If we let $\mathcal{P} = \{I, X, Y, Z\}$ be the Pauli matrices, then we can define a generalization of the Pauli matrices to an n -qubit system as $\mathcal{P}_n = \{\pm 1\} \cdot \{I, X, Y, Z\}^{\otimes n}$.³ As an example, $-X \otimes Y \otimes Z \otimes Y \otimes Z \in \mathcal{P}_5$, though in the future we will simply opt to write this Pauli matrix as $-XYZYZ$. It is not hard to show that $\forall P \neq \pm I^{\otimes n} \in \mathcal{P}_n$ half of the eigenvalues are 1 and the other half are -1 .

We say that a pure state $\rho = |\psi\rangle\langle\psi|$ such that $|\psi\rangle \in \mathbb{C}^{2^n}$ is *stabilized* by $P \in \mathcal{P}_n$ if $P|\psi\rangle = |\psi\rangle$. In other words, $|\psi\rangle$ must be an eigenvector of P with eigenvalue 1. The set of pure states that are stabilized by the subset $S \subseteq \mathcal{P}_n$ are then the states that lie in the intersection of the eigenvalue 1 subspaces. It is known that if S is an Abelian group that does not contain $-I^{\otimes n}$, then this intersection has dimension $2^n/|S|$ [29]. Due to the nature of the outer product, any vector $e^{i\theta}|\psi\rangle$ drawn from this space results in the same density matrix $\rho = |\psi\rangle\langle\psi|$.

Definition 4.1 (Stabilizer states). Let $S \subset \mathcal{P}_n \setminus \{-I^{\otimes n}\}$ be an Abelian group of order 2^n (note that \mathcal{P}_n is not itself a group under matrix multiplication). The unique state density matrix $\rho = |\psi\rangle\langle\psi|$ that results from the one-dimensional subspace that is stabilized by all elements of S is then defined to be a stabilizer state. We then say that S stabilizes ρ .

The stabilizer states are then the set of all such quantum pure states that are stabilized by Abelian groups of order 2^n formed from $\mathcal{P}_n \setminus \{-I^{\otimes n}\}$.

Proposition 4.2. *Given two n -qubit stabilizer states $|\psi\rangle\langle\psi| \neq |\varphi\rangle\langle\varphi|$ with stabilizer groups S and S' respectively, then $|S \cap S'| \leq 2^{n-1}$.*

Proof. Because $|\psi\rangle\langle\psi| \neq |\varphi\rangle\langle\varphi|$ then $S \neq S'$. We also know that $S \cap S'$ is an abelian group without $-I^{\otimes n}$, so $|S \cap S'| < 2^n$. Since $2^n/|S \cap S'|$ is the dimension of the space stabilized by this group [29], it must be an integer. Due to the prime factorization of 2^n , $|S \cap S'| = 2^m$ for some integer $0 \leq m < n$, of which the largest possible m is $n - 1$. \square

4.2 Difficulty of beating the maximally mixed state on uniform Pauli measurements

Let \mathcal{C} be the class of all n -qubit stabilizer pure states. If $P \in \mathcal{P}_n$ is a Pauli operator, then the two-outcome measurement associated with P is $(I + P)/2$, and is referred to as a Pauli mea-

³Note that this is slightly different from the usual Pauli group $\{\pm 1, \pm i\} \cdot \{I, X, Y, Z\}^{\otimes n}$. We choose to ignore the imaginary phases as they will never be a part of a valid stabilizer measurement: for instance, $(I + iX)/2$ is not even Hermitian.

surement. We will first examine the natural distribution D given by the uniform distribution over Pauli measurements,

$$\mathcal{E}_{\text{Pauli}} = \{(I + P)/2 \mid P \in \mathcal{P}_n\}.$$

In doing so, we will show that performing better than the trivial algorithm of always outputting the maximally mixed state $I/2^n$ is difficult.

We use the following folklore lemma (see e.g. [41] Lemma 1 for a simple proof).

Lemma 4.3. *Let $E = (I + P)/2$ be a POVM measurement associated to a Pauli operator $P \in \mathcal{P}_n$ and ρ be an n -qubit stabiliser state. Then $\text{Tr}(E\rho)$ can only take on the values $\{0, \frac{1}{2}, 1\}$, and:*

$$\begin{cases} \text{Tr}(E\rho) = 1 \text{ iff } P \text{ is a stabilizer of } \rho; \\ \text{Tr}(E\rho) = 1/2 \text{ iff neither } P \text{ nor } -P \text{ is a stabilizer of } \rho; \\ \text{Tr}(E\rho) = 0 \text{ iff } -P \text{ is a stabilizer of } \rho. \end{cases}$$

Simple algebraic manipulations then tell us that $\text{Tr}(P\rho)$ can only take on the values $\{1, 0, -1\}$ with $\text{Tr}(P\rho)$ being 1 or -1 if and only if P or $-P$ is in the stabilizer group of ρ respectively. With this result, we can compute bounds on $|\langle f_\rho, f_{\rho'} \rangle_D|$ for stabilizer states ρ and ρ' by counting how many matrices lie in the intersection of their stabilizer groups or the negations of their stabilizer groups.

Lemma 4.4. *Let \mathcal{C} be the concept class of n -qubit stabilizer pure states, and let D denote the uniform distribution on n -qubit Pauli measurements. Then for any stabilizer states ρ, ρ' with $\rho \neq \rho'$, $|\langle f_\rho, f_{\rho'} \rangle_D| = \|f_\rho\|_D^2 = \frac{1}{2^n}$, and $|\langle f_\rho, f_{\rho'} \rangle_D| \leq \frac{1}{2^{n+1}}$. Furthermore, this inequality is tight.*

Proof. Let $\rho, \rho' \in \mathcal{C}$. Let S and S' be the stabilizer groups for ρ and ρ' respectively, and also let $-S = \{-P \mid P \in S\}$ and $-S' = \{-P \mid P \in S'\}$. For any Pauli measurement $(I + P)/2$, we have $f_\rho(\frac{I+P}{2}) = 2 \text{Tr}(\frac{I+P}{2}\rho) - 1 = \text{Tr}(P\rho)$. Thus the correlation of the two p-concepts becomes

$$|\langle f_\rho, f_{\rho'} \rangle_D| = \frac{1}{|\mathcal{P}_n|} \left| \sum_{P \in \mathcal{P}_n} \text{Tr}(P\rho) \cdot \text{Tr}(P\rho') \right|.$$

Leveraging Lemma 4.3, we find:

$$\begin{aligned} |\langle f_\rho, f_{\rho'} \rangle_D| &= \frac{1}{|\mathcal{P}_n|} \left| \sum_{P \in \mathcal{P}_n} \text{Tr}(P\rho) \cdot \text{Tr}(P\rho') \right| \\ &= \frac{1}{2 \cdot 4^n} [|S \cap S'| + | -S \cap -S'| - |S \cap -S'| - | -S \cap S'|] \\ &= \frac{1}{4^n} [|S \cap S'| - |S \cap (-S')|] \end{aligned}$$

If $\rho = \rho'$ then $S = S'$ such that $|S \cap S'| = 2^n$ and $|S \cap -S'| = 0$. Thus for all ρ

$$|\langle f_\rho, f_\rho \rangle_D| = \frac{2^n}{4^n} = \frac{1}{2^n}$$

If $\rho \neq \rho'$ then by Proposition 4.2 $|S \cap S'| \leq 2^{n-1}$. $|S \cap -S'| \geq 0$ trivially, so we get an upper bound of

$$|\langle f_\rho, f_{\rho'} \rangle_D| \leq \frac{1}{4^n} (2^{n-1}) = \frac{1}{2^{n+1}}.$$

We can show that this inequality is tight because the state $\rho = |0\rangle\langle 0|^{\otimes n}$ and $\rho' = |0\rangle\langle 0|^{\otimes n-1} \otimes |+\rangle\langle +|$ saturate this inequality. The generators for the stabilizer group of ρ are (omitting tensor products): $ZIII \cdots I$, $IZIII \cdots I$, \dots , and $IIII \cdots IZ$. The generators of ρ' are the same, except the last generator is replaced with $IIII \cdots IX$. We see that $|S \cap S'| = 2^{n-1}$ while $|S \cap -S'| = 0$.

□

With this result, we can use Lemma 2.5 to compute the SDA and by extension prove a lower bound on the number of statistical queries needed to learn this concept class under this distribution.

Proposition 4.5 ([3] Proposition 2). *The number of n -qubit stabilizer states grows as $2^{\Theta(n^2)}$.*

Theorem 4.6. *Let D be the uniform distribution over n -qubit Pauli measurements and let \mathcal{C} be the concept class of all n -qubit stabilizer pure states. Then $\text{SDA}(\mathcal{C}, \frac{1}{2^n}) = 2^{\Theta(n^2)}$.*

Proof. By Proposition 4.5, $|\mathcal{C}| = 2^{\Theta(n^2)}$. Using Lemma 2.5 with $\kappa = \frac{1}{2^n}$ and $\gamma = \frac{1}{2^{n+1}}$ as calculated from Lemma 4.4, we find that

$$\text{SDA}(\mathcal{C}, \gamma' + \frac{1}{2^{n+1}}) = \text{SDA}(\mathcal{C}, \gamma' + \gamma) \geq |\mathcal{C}| \frac{\gamma'}{\beta - \gamma} = 2^{\Theta(n^2)} \gamma' 2^{n+1} = 2^{\Theta(n^2)} \gamma'$$

Setting $\gamma' = \frac{1}{2^{n+1}}$ gives the result.

□

Corollary 4.7. *Any SQ algorithm needs at least $2^{\Omega(n^2)}$ statistical queries of tolerance $\tau = 2^{-O(n)}$ to learn \mathcal{C} up to error $2^{-O(n)}$ over D .*

Proof. Simply apply Theorem 2.4, with $\beta = 2^{-n}$.

□

Since the norms of our p-concepts are exponentially small (i.e. $2^{-n/2}$), we only get hardness for error on the order of $2^{-O(n)}$. But as we now show, the p-concept norm corresponds almost exactly to the squared loss achieved by the maximally mixed state. Our results show that doing significantly better than the maximally mixed state requires $2^{\Omega(n^2)}$ statistical queries even when the tolerance is exponentially small.

Proposition 4.8. *Let D be the uniform distribution over n -qubit Pauli measurements, $\mathcal{E}_{\text{Pauli}}$. Let ρ be any state, and let $I/2^n$ be the maximally mixed state. Then*

$$\|f_\rho\|_D^2 = \|f_\rho - f_{I/2^n}\|_D^2 + 1/4^n.$$

Proof. In essence, this is simply because the p-concept $f_{I/2^n}$ is almost always zero. Specifically, for all $E \in \mathcal{E}_{\text{Pauli}} \setminus \{0, I\}$, $f_{I/2^n}(E) = 2 \text{Tr}(E/2^n) - 1 = 0$, since $\text{Tr}(E) = 2^{n-1}$ for all such E . (This is because $E = (I + P)/2$ for some n -qubit Pauli matrix $P \in \mathcal{P}_n$, and $\text{Tr}(P) = 0$ for all $P \in \mathcal{P}_n \setminus \{\pm I\}$.) As for $E \in \{0, I\}$, we note that $f_\rho(E) = f_{I/2^n}(E)$. Thus

$$\begin{aligned}
\|f_\rho\|_D^2 &= \frac{1}{|\mathcal{E}_{\text{Pauli}}|} \sum_{E \in \mathcal{E}_{\text{Pauli}}} f_\rho(E)^2 \\
&= \frac{1}{|\mathcal{E}_{\text{Pauli}}|} \left(\sum_{E \in \mathcal{E}_{\text{Pauli}} \setminus \{0, I\}} f_\rho(E)^2 + \sum_{E \in \{0, I\}} f_\rho(E)^2 \right) \\
&= \frac{1}{|\mathcal{E}_{\text{Pauli}}|} \left(\sum_{E \in \mathcal{E}_{\text{Pauli}} \setminus \{0, I\}} (f_\rho(E) - f_{I/2^n}(E))^2 + \sum_{E \in \{0, I\}} f_\rho(E)^2 \right) \\
&= \frac{1}{|\mathcal{E}_{\text{Pauli}}|} \left(\sum_{E \in \mathcal{E}_{\text{Pauli}}} (f_\rho(E) - f_{I/2^n}(E))^2 + \sum_{E \in \{0, I\}} f_\rho(E)^2 \right) \\
&= \|f_\rho - f_{I/2^n}\|_D^2 + \frac{2}{|\mathcal{E}_{\text{Pauli}}|} \\
&= \|f_\rho - f_{I/2^n}\|_D^2 + 1/4^n.
\end{aligned}$$

□

4.3 Lower bounds via a direct reduction from learning parities

To get around this norm issue, we look at a subset of stabilizer states such that we can produce p-concepts with norm 1. Recall that the Pauli measurements are the set of all projectors onto the eigenvalue-1 space of some Pauli matrix P , i.e. $\{\frac{P+I}{2} \mid P \in \mathcal{P}_n\}$. We define a subset of the Pauli measurements called the parity measurements, and show the hardness of SQ-learning stabilizer states under the uniform distribution on such measurements. This is via a simple equivalence, holding essentially by construction, with the problem of learning parities under the uniform distribution. As a further consequence, we obtain that learning stabilizer states with noise is at least as hard as Learning Parity with Noise (LPN). This holds for general PAC-learning, even outside the SQ model.

Definition 4.9 (Parity measurements). For all $x \in \{0, 1\}^n$, let $P_x = \sum_{y \in \{0, 1\}^n} \chi_x(y) |y\rangle\langle y|$. Since the set of P_x is equivalent to $\{I, Z\}^{\otimes n}$, the corresponding measurement $E_x = \frac{P_x + I}{2}$ is by definition a Pauli measurement. We will refer to such measurements as *parity measurements*.

Proposition 4.10. *For every distribution D on $\{0, 1\}^n$ there is a corresponding distribution D' on parity measurements such that learning computational basis states under D' is equivalent to learning parities under D . Furthermore, this equivalence holds even with classification noise: for any η , learning computational basis states under D' with noise rate η is equivalent to learning parities under D with noise rate η .*

In particular, learning stabilizer states under D' is at least as hard as learning parities under D .

Proof. If the unknown state ρ is a computational basis state $|y\rangle\langle y|$, then the value

$$\text{Tr}(E_x |y\rangle\langle y|) = x \cdot y \pmod{2}$$

is simply the parity of x over the subset specified by y (represented using $\{0, 1\}$ instead of $\{-1, 1\}$). In the PAC setting, this would be equivalent to getting the sample $(E_x, x \cdot y \pmod{2})$.

Accordingly, let us define D' simply as the distribution over E_x for $x \sim D$. It is clear that these are different representations of the same problem, such that a learning algorithm for one implies a learning algorithm for the other. We note that this relationship holds even in the presence of classification noise. Finally, note that computational basis states are a subset of the stabilizer states, so any learner for stabilizer states implies a learner for the computational basis states as well. This implies that learning stabilizer states on D' is at least as hard as learning parities on D , even in the presence of classification noise. \square

Corollary 4.11. *SQ-learning stabilizer states under the uniform distribution over parity measurements requires $2^{\Omega(n)}$ queries even with constant error (say $1/3$).*

Proof. By Proposition 4.10, SQ-learning stabilizer states under the uniform distribution on E_x parity measurements is at least as hard as learning parities over the uniform distribution. Applying Theorem 2.6, we get the exponential lower bound. \square

Corollary 4.12. *Learning stabilizer states under the uniform distribution over parity measurements with classification noise rate η is at least as hard as LPN with noise rate η .*

Proof. Proposition 4.10 directly implies that learning computational basis states under the uniform distribution on parity measurements and with classification noise is equivalent to LPN. \square

5 An SQ learner for product states

Turning to positive results, we now give SQ algorithms for some simple concept classes, namely the computational basis states and, more generally, products of n single-qubit states. The distribution on measurements that we will consider will correspond to a natural scheme for these classes: pick a qubit at random and measure it using a Haar-random unitary.

Concretely, let D' be the distribution of single qubit measurements formed from the projection onto Haar-random single qubit state (i.e. $U|0\rangle\langle 0|U^\dagger$ where U is a Haar random unitary), and let D be the distribution on n -qubit measurements that corresponds to picking a qubit at random and measuring it using a measurement drawn from D' . That is, $D = \frac{1}{n} \sum_{i=1}^n I^{\otimes i-1} \otimes D' \otimes I^{\otimes n-i}$. Let \mathcal{C} be the concept class of product states $\rho = \otimes_{i=1}^n \rho_i$. Of course, this class includes the computational basis states. The main result of this section will be a simple $O(n)$ -query SQ algorithm for learning \mathcal{C} under the distribution D .

We remark that our algorithm's guarantee actually trivially extends to learning arbitrary (not just product) states under the above distribution D of single-qubit Haar-random measurements. This is simply because such measurements only ever inspect each qubit individually, so that a product state $\otimes_i \rho_i$ is indistinguishable—*under D* —from a more general mixed state ρ whose reduced density matrix on qubit i is ρ_i for every i .⁴ Yet since this distribution on measurements is fundamentally not very interesting for anything other than product states, we state the results in this section only for product states.

The following technical lemma will be the backbone of our results.

⁴We stress that in the PAC formalism, the goal is not necessarily to learn the exact state, but simply to find one that behaves similarly under the specified input distribution of measurements. Thus for measurements of the kind drawn from D , learning the product of reduced density states is sufficient.

Lemma 5.1. For any single qubit pure state $|\psi\rangle\langle\psi| = \frac{I+P}{2}$ and mixed state ρ :

$$\mathbb{E}_{E \sim D'} \left[\text{sign} \left(\text{Tr}(E|\psi\rangle\langle\psi|) - \frac{1}{2} \right) \left(\text{Tr}(E\rho) - \frac{1}{2} \right) \right] = \frac{1}{4} \text{Tr}(P\rho).$$

Proof. We will decompose $\rho = \lambda|\varphi\rangle\langle\varphi| + (1-\lambda)|\varphi^\perp\rangle\langle\varphi^\perp|$ such that $|\varphi\rangle = \cos\theta'|\psi\rangle + \sin\theta'|\psi^\perp\rangle$ and $|\varphi^\perp\rangle = e^{i\varphi'}(\sin\theta'|\psi\rangle - \cos\theta'|\psi^\perp\rangle)$. The following identity will be useful at the end:

$$\begin{aligned} \text{Tr}(P\rho) &= 2 \text{Tr}(|\psi\rangle\langle\psi|\rho) - 1 \\ &= 2[\lambda \text{Tr}(|\psi\rangle\langle\psi||\varphi\rangle\langle\varphi|) + (1-\lambda) \text{Tr}(|\psi\rangle\langle\psi||\varphi^\perp\rangle\langle\varphi^\perp|)] - 1 \\ &= 2 \cos^2\theta'\lambda + 2(1-\lambda) \sin^2\theta' - (\sin^2\theta' + \cos^2\theta') \\ &= (2\lambda - 1) \cos^2\theta' - (2\lambda - 1) \sin^2\theta' \\ &= (2\lambda - 1) \cos 2\theta' \end{aligned}$$

Let U be the unitary such that $U|0\rangle = |\psi\rangle$ and $U|1\rangle = |\psi^\perp\rangle$. Due to symmetry, we can parameterize a Haar-random single qubit state using spherical coordinates as $E = \frac{1}{2}U(I + \cos\varphi \sin\theta X + \sin\varphi \sin\theta Y + \cos\theta Z)U^\dagger$ for the Pauli matrices X, Y , and Z .

$$\begin{aligned} \text{Tr}(E|\psi\rangle\langle\psi|) &= \text{Tr} \left(\frac{1}{2}U(I + \cos\varphi \sin\theta X + \sin\varphi \sin\theta Y + \cos\theta Z)U^\dagger|\psi\rangle\langle\psi| \right) \\ &= \text{Tr} \left(\frac{1}{2}(I + \cos\varphi \sin\theta X + \sin\varphi \sin\theta Y + \cos\theta Z)|0\rangle\langle 0| \right) \\ &= \frac{1 + \cos\theta}{2} \end{aligned}$$

We can also do the same thing for ρ :

$$\begin{aligned} \text{Tr}(E\rho) &= \lambda \text{Tr}(E|\varphi\rangle\langle\varphi|) + (1-\lambda) \text{Tr}(E|\varphi^\perp\rangle\langle\varphi^\perp|) \\ &= \lambda \frac{1 + \cos\theta \cos 2\theta' + \cos\varphi \sin\theta \sin 2\theta'}{2} + (1-\lambda) \frac{1 - \cos\theta \cos 2\theta' - \cos\varphi \sin\theta \sin 2\theta'}{2} \\ &= \frac{1 + (2\lambda - 1)(\cos\theta \cos 2\theta' + \cos\varphi \sin\theta \sin 2\theta')}{2} \end{aligned}$$

This allows us to perform a spherical integral over θ and φ to get the expectation:

$$\begin{aligned}
& \mathbb{E}_{E \sim D'} \left[\text{sign} \left(\text{Tr}(E|\psi\rangle\langle\psi|) - \frac{1}{2} \right) \left(\text{Tr}(E|\varphi\rangle\langle\varphi|) - \frac{1}{2} \right) \right] \\
&= \frac{1}{4\pi} \int_0^{2\pi} d\varphi \int_0^\pi d\theta \sin \theta \text{sign}(\cos \theta) \left((2\lambda - 1) \frac{\cos \theta \cos 2\theta' + \cos \varphi \sin \theta \sin 2\theta'}{2} \right) \\
&= \frac{2\lambda - 1}{8\pi} \int_0^{2\pi} d\varphi \left[\int_0^{\pi/2} - \int_{\pi/2}^\pi \right] d\theta \sin \theta (\cos \theta \cos 2\theta' + \cos \varphi \sin \theta \sin 2\theta') \\
&= (2\lambda - 1) \frac{\pi \cos 2\theta' + \pi \cos 2\theta'}{8\pi} \\
&= (2\lambda - 1) \frac{\cos 2\theta'}{4} \\
&= \frac{1}{4} \text{Tr}(P\rho)
\end{aligned}$$

□

Our algorithm for learning product states will be work by learning each qubit in the Pauli basis. This results in a $3n$ -query algorithm, corresponding to the $3n$ parameters that it takes to define a product state. We first recall the definition of trace distance, which is the quantum generalization of total variational distance.

Definition 5.2 (Trace distance). Given quantum states ρ and σ ,

$$\|\rho - \sigma\|_{\text{Tr}} = \frac{1}{2} \sum_i |\lambda_i|$$

where $\{\lambda_i\}$ is the set of eigenvalues of the matrix $\rho - \sigma$.

Proposition 5.3 (folklore). *Given two single qubit states ρ and σ , the $\|\rho - \sigma\|_{\text{Tr}}$ is half the Euclidean distance between their points on the Bloch sphere.*

The following lemma will then be necessary to relate trace distance of the states to the squared loss in learning under this distribution.

Lemma 5.4. *For n -qubit product states $\rho = \bigotimes_i \rho_i$ and $\sigma = \bigotimes_i \sigma_i$, let $f_\rho(E) = 2 \text{Tr}(E\rho) - 1$ and $f_\sigma(E) = 2 \text{Tr}(E\sigma) - 1$. Let D be the distribution over measurements defined earlier. Then*

$$\mathbb{E}_{E \sim D} [(f_\rho(E) - f_\sigma(E))^2] = \frac{4}{3n} \sum_{i=1}^n \|\rho_i - \sigma_i\|_{\text{Tr}}^2$$

Proof. Let $\xi = \rho - \sigma$. Then by linearity

$$f_\rho(E) - f_\sigma(E) = 2(\text{Tr}(E\rho) - \text{Tr}(E\sigma)) = 2 \text{Tr}(E\xi).$$

We will define $\xi_i = \text{Tr}_i(\xi) = \rho_i - \sigma_i$ to be the reduced density matrix on the i^{th} qubit of ξ . Noting that each ξ_i is traceless, then by diagonalizing we can write $\xi_i = \lambda_i |\lambda_i\rangle\langle\lambda_i| - \lambda_i |\lambda_i^\perp\rangle\langle\lambda_i^\perp|$ for $\lambda_i \in [0, 1]$ such that $\lambda_i = \|\rho_i - \sigma_i\|_{\text{Tr}}$ is the trace distance of the reduced density matrix.

Like in Lemma 5.1, we can parameterize a single-qubit Haar-random projection as $E = \frac{1}{2}U(I + \cos \varphi \sin \theta X + \sin \varphi \sin \theta Y + \cos \theta Z)U^\dagger$, where $U|0\rangle = U|\lambda_i\rangle$ and $U|1\rangle = U|\lambda_i^\perp\rangle$. This implies that $U\xi_i U^\dagger = \lambda_i Z$.

$$\begin{aligned}\mathrm{Tr}(E\xi_i) &= \mathrm{Tr}\left(\frac{1}{2}U(I + \cos \varphi \sin \theta X + \sin \varphi \sin \theta Y + \cos \theta Z)U^\dagger \rho\right) \\ &= \mathrm{Tr}\left(\frac{1}{2}(I + \cos \varphi \sin \theta X + \sin \varphi \sin \theta Y + \cos \theta Z) \cdot \lambda_i Z\right) \\ &= \lambda_i \cos \theta\end{aligned}$$

Using this, we now compute the squared-loss as follow.

$$\begin{aligned}\mathbb{E}_{E \sim D} [(f_\rho(E) - f_\sigma(E))^2] &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{E' \sim D'} [(f_{\rho_i}(E') - f_{\sigma_i}(E'))^2] \\ &= \frac{4}{n} \sum_{i=1}^n \mathbb{E}_{E' \sim D'} [\mathrm{Tr}^2(E'\xi_i)] \\ &= \frac{4}{n} \sum_{i=1}^n \frac{1}{2} \int_0^\pi \sin \theta \cdot \lambda_i^2 \cos^2 \theta \\ &= \frac{4}{3n} \sum_{i=1}^n \lambda_i^2 \\ &= \frac{4}{3n} \sum_{i=1}^n \|\rho_i - \sigma_i\|_{\mathrm{Tr}}^2\end{aligned}$$

□

We now show how to use Lemma 5.1 to learn each qubit of the product state, allowing us to then apply Lemma 5.4 to get our learning result.

Theorem 5.5. *Let D be the distribution on measurements and let \mathcal{C} be the concept class of product states defined earlier. There is an SQ learner that is able to learn \mathcal{C} under D up to squared loss ϵ using $3n$ queries of tolerance $\sqrt{\epsilon/n}$.*

Proof. Let the unknown $\rho \in \mathcal{C}$ be given by $\rho = \bigotimes_i \rho_i$. If we define $P_1 = X$, $P_2 = Y$, and $P_3 = Z$, then our queries will be

$$\varphi_{i,j}(E, Y) = \mathrm{sign}\left(\frac{1}{2^{n-1}} \mathrm{Tr}\left(E \cdot \left(I^{\otimes i-1} \otimes \frac{I + P_j}{2} \otimes I^{\otimes n-i}\right)\right) - \frac{1}{2}\right) \cdot Y$$

The query $\varphi_{i,j}$ will correspond to taking the projection of the i^{th} qubit along the Pauli P_j , as

we now show:

$$\begin{aligned}
\mathbb{E}_{E \sim D, Y \sim f_\rho(E)} [\varphi_{i,j}(E, Y)] &= \mathbb{E}_{E \sim D} \left[\varphi_{i,j}(E, 1) \text{Tr}(E\rho) + \varphi_{i,j}(E, -1) \left(1 - \text{Tr}(E\rho)\right) \right] \\
&= \mathbb{E}_{E \sim D} \left[\varphi_{i,j}(E, 1) \left(2 \text{Tr}(E\rho) - 1\right) \right] \\
&= \frac{1}{n} \mathbb{E}_{E' \sim D'} \left[\text{sign} \left(\text{Tr}(E'|0\rangle\langle 0|) - \frac{1}{2} \right) \left(2 \text{Tr}(E'\rho_i) - 1\right) \right] \\
&= \frac{1}{2n} \text{Tr}(P_j \rho_i).
\end{aligned}$$

Here the third equality exploits the definition of D as $\frac{1}{n} \sum_{i=1}^n I^{\otimes i-1} \otimes D' \otimes I^{\otimes n-i}$ (only the i^{th} term yields a nonzero expectation), and the fourth equality is Lemma 5.1.

Any specific qubit ρ_i can be written in Bloch sphere coordinates as $\frac{I+x_iX+y_iY+z_iZ}{2}$. We can estimate $x_i = \frac{1}{2} \text{Tr}(P_1 \rho_i)$ up to error $\sqrt{\epsilon}$ using a single query of tolerance $\sqrt{\epsilon}/n$. The same holds true for y_i and z_i . If we use this to construct our estimate

$$\hat{\rho}_i = \frac{I + \hat{x}_i X + \hat{y}_i Y + \hat{z}_i Z}{2}$$

then by Proposition 5.3 we get

$$\|\rho_i - \hat{\rho}_i\|_{\text{Tr}}^2 = \frac{1}{4} \left[(x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2 + (z_i - \hat{z}_i)^2 \right] \leq 3\epsilon/4.$$

Finally, using Lemma 5.4:

$$\mathbb{E}_{E \sim D} [(f_\rho(E) - f_\sigma(E))^2] \leq \frac{4}{3n} \sum_{i=1}^n 3\epsilon/4 = \epsilon.$$

We note that if the estimated point is outside of the Bloch sphere, we can simply normalize the point to the surface of the Bloch sphere and this will never increase the error. To quickly sketch the proof of this, take the plane formed by the center of the sphere, the estimated point \hat{p} that is outside of the sphere, and the real point p which is both within the Bloch sphere and within a sphere ϵ radius located at \hat{p} . The normalized point \hat{p}' is always located on the line from the \hat{p} to the origin, and one can make a separating plane that bisects the line segment between \hat{p} and \hat{p}' that denotes whether one is closer to \hat{p} or \hat{p}' . Since the Bloch sphere will always be on the side closer to \hat{p}' and the real point p is in the Bloch sphere, p will always be closer to \hat{p}' than \hat{p} . \square

We can simplify this algorithm if we know in advance that ρ is a computational basis state. In that case, we know that each qubit ρ_i is either $(I+Z)/2$ or $(I-Z)/2$, and so we only need to make n queries $\varphi_{i,3}$, one for each i . Moreover, we only need to identify the coordinate z_i to within an accuracy of 1 in order to distinguish the $z_i = 1$ and $z_i = -1$ cases, so that our tolerance need only scale as $O(1/n)$ in order to learn ρ perfectly (i.e. with $\epsilon = 0$).

6 Connections to differential privacy

A PAC learning algorithm L can be viewed as a randomized algorithm that takes as input a training dataset (i.e. a set of labeled examples (x, y) sampled from a distribution) and outputs

a hypothesis that with high probability has low error over the distribution. That is, if S is a training dataset, then $L(S)$ describes a probability distribution over hypotheses (where the randomness arises from the internal randomness of the learner). Intuitively, differential privacy requires L to satisfy a kind of stability: on any two inputs S and S' that are close, the distributions $L(S)$ and $L(S')$ must be close as well.

Definition 6.1 (Differential privacy, [21]). Call two datasets $S = \{(x_i, y_i)\}_{i=1}^m$ and $S' = \{(x'_i, y'_i)\}_{i=1}^m$ neighbors if they only differ in one entry. A learner L (understood in the sense just discussed) is said to be α -differentially private (or α -DP for short) if for any S and S' that are neighbors, the distributions $L(S)$ and $L(S')$ are close in the sense that for any hypothesis h , $\mathbb{P}[L(S) = h] \leq e^\alpha \mathbb{P}[L(S') = h]$.

A well-known property of SQ algorithms is that they can readily be made differentially private [17, 21]. Since differential privacy is a notion that is well-defined only in the PAC setting where the input is a set of training examples (as opposed to access to an SQ oracle), such a statement is necessarily of the form “any SQ learner yields a PAC learner that satisfies differential privacy.”

Theorem 6.2 (see e.g. [13]). *Let \mathcal{C} be a concept class learnable up to error ϵ by an SQ learner L using q queries of tolerance τ . Then it is also learnable up to error ϵ in the PAC setting by an α -DP learner L' with sample complexity $\tilde{O}(\frac{q}{\alpha\tau} + \frac{q}{\tau^2})$ (with constant probability).*

The proof is standard and proceeds by simulating each of L 's queries using empirical estimates over a sample of size roughly $1/\tau^2$ and then using the Laplace mechanism to add some further noise.

One can extend this notion to the quantum setting. One natural and direct way of doing so is simply by replacing the classical dataset of labeled pairs (x_i, y_i) by one of measurement-outcome pairs (E_i, Y_i) ; the rest remains exactly analogous. Theorem 6.2 then carries over verbatim to our notion of quantum SQ learnability. This form of quantum differential privacy was recently studied by Arunachalam et al. [11], who were able to relate it to online learning, one-way communication complexity, and shadow tomography of quantum states, extending ideas of Bun et al. [18]. Since our notion of quantum SQ learnability implies quantum DP learnability, it also fits into their framework. In particular, by the chain of implications established in that work, efficient quantum SQ learnability of a class of states implies DP PAC learnability, which implies finite sequential fat-shattering (sfat) dimension, which in turn implies online learnability, gentle shadow tomography, and “quantum stability.” In fact, in the classical setting, some of the main examples of realistic DP learners are SQ (even though technically the inclusion is known to be strict) [17, 31], and one might expect the same to hold in the quantum setting as well.

We remark that a somewhat different kind of quantum differential privacy, where privacy is with respect to copies of the unknown state, may also be defined as follows. View a quantum state learner L as an algorithm that takes in multiple copies $\rho^{\otimes m}$ of some unknown state ρ , is allowed to sample and perform random measurements from a distribution D , and outputs another state σ that is required to be close to ρ with respect to D with high probability. If the random measurements are viewed as the internal randomness of the learner, then this is similar to the view we took of a classical learner earlier. We can now define a notion of differential privacy for quantum state learners by requiring that $L(\rho^{\otimes m})$ and $L(\rho^{\otimes m-1} \otimes \rho')$

(where $\rho \neq \rho'$, so that $\rho^{\otimes m}$ and $\rho^{\otimes m-1} \otimes \rho'$ are neighbors) are α -close as distributions over states (in the natural way). This can also be seen as a stylized kind of tolerance to noise or corruptions. The following analogue of Theorem 6.2 can then be proven using almost exactly the same proof; essentially, we are only replacing classical examples with copies of quantum states.

Theorem 6.3. *Let \mathcal{C} be a class of quantum states learnable up to error ϵ by an SQ learner L using q queries of tolerance τ . Then it is also learnable up to error ϵ in the PAC setting by an α -DP learner L' (in the specific sense just described) with copy complexity $\tilde{O}(\frac{q}{\alpha\tau} + \frac{q}{\tau^2})$ (with constant probability).*

Note that these notions are different from those of [5], which defined differential privacy for quantum measurements. Here two n -qubit states are considered neighbors if it is possible to reach one from the other by a quantum operation (sometimes called a superoperator) on a single qubit. In particular, two product states $\rho = \otimes_i \rho_i$ and $\sigma = \otimes_i \sigma_i$ are neighbors if $\rho_i = \sigma_i$ for all i but one.

Definition 6.4 (Quantum differential privacy for measurements, [5]). A measurement M is said to be α -DP if for any n -qubit neighbor states ρ, σ , and any outcome y , $\mathbb{P}[M(\rho) = y] \leq e^\alpha \mathbb{P}[M(\sigma) = y]$.

The authors show that this definition can be related to the notion of a “gentle quantum measurement,” and this connection can be carefully exploited to perform shadow tomography [2]. However, this kind of quantum DP is not applicable in a natural way to a PAC or SQ learner, since such a learner is an algorithm rather than just a single measurement.

7 Discussion and open problems

Statistical vs. query complexity. Conceptually, the contrast between our SQ model and the original PAC model of [1] is interesting. Apart from the definition of an elegant model, Aaronson’s main insight was to characterize learnability in a purely *statistical* sense, showing bounds on sample complexity via an analysis of the so-called fat-shattering dimension of quantum states. In learning theoretic terms, this took advantage of a separation of concerns that the PAC model encourages: (a) empirical performance, i.e. a learner achieving low error with respect to the training data, and (b) generalization, i.e. this performance actually generalizing to the true distribution. The SQ model, however, does not naturally accommodate such a separation. SQ algorithms are instead primarily characterized by the number of queries required; generalization is “in-built.” The closest analogue to a notion of sample complexity is the role played by the tolerance, and the closest thing to studying generalization on its own might have been to show a phase transition in what different regimes of the tolerance are able to accomplish. The formal statements of our SQ lower bounds do have such a flavor: “*either use small tolerance or many queries.*”

Suitable classes and distributions for PAC-learning. It is notable that the algorithms of [41] for learning stabilizer states and [45] for low Schmidt rank states are essentially the only known positive results in the framework of [1]. Both these algorithms are “algebraic” and involve solving a system of polynomial equations, something that SQ cannot handle. A

longstanding question in this area is: what other interesting classes can be learned, and under what distributions on measurements? And can they also be learned in the SQ setting?

A major issue in picking suitable distributions on measurements is that under many natural distributions, the maximally mixed state actually performs quite well, so that the problem of learning becomes essentially superfluous. Even in this work, we obtained lower bounds for learning stabilizer states under the uniform distribution on Pauli measurements only for learning up to exponentially small squared loss. This was because the norms of the p-concepts are themselves exponentially small, or in other words the maximally mixed state already achieves exponentially small loss. We were able to get around this and obtain a $\Omega(2^n)$ lower bound via a direct reduction from learning parities (by considering parity measurements). Can we do better than just 2^n ? Is there a $\omega(2^n)$ -sized (e.g., 4^n or 2^{n^2}) subset of stabilizer states such that there exists a distribution over Pauli measurements inducing norms that are only polynomially small yet have an exponentially small average correlation? That is, is there a $\omega(2^n)$ -sized set of stabilizer states and accompanying distribution over Pauli measurements such that the maximally mixed state does not do well?

Other forms of noise. Can we extend the noise tolerance of SQ algorithms to more forms of noise, or improve the parameters of the noise tolerated? One such interesting form of noise would be depolarizing noise that acts on individual qubits (as opposed to acting directly on the whole state).

Noise-tolerant learning beyond SQ. The best-known PAC algorithm for learning parities with noise is due to [16] and runs in slightly subexponential time. Interestingly, this means it beats the exponential SQ lower bound and is hence essentially the only known example of a noise-tolerant PAC algorithm that is not SQ (although it cannot handle noise arbitrarily close to the information-theoretic limit). Can we similarly hope for a noise-tolerant but non-SQ learner for stabilizer states that runs in subexponential time?

Acknowledgements

We thank Scott Aaronson, Srinivasan Arunachalam, and Andrea Rocchetto for many helpful discussions. AG was supported by NSF awards AF-1909204, AF-1717896, and the NSF AI Institute for Foundations of Machine Learning (IFML). DL was supported by the Simons It from Qubit Collaboration and Scott Aaronson’s Vannevar Bush Faculty Fellowship from the US Department of Defense.

References

- [1] Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, 2007. DOI: [10.1098/rspa.2007.0113](https://doi.org/10.1098/rspa.2007.0113). 1, 2, 3, 4, 6, 21
- [2] Scott Aaronson. Shadow tomography of quantum states. *SIAM Journal on Computing*, 49(5):STOC18–368, 2019. DOI: [10.1145/3188745.3188802](https://doi.org/10.1145/3188745.3188802). 21
- [3] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004. DOI: [10.1103/physreva.70.052328](https://doi.org/10.1103/physreva.70.052328). 2, 13

- [4] Scott Aaronson and Sabee Grewal. Efficient learning of non-interacting fermion distributions, 2021. 4
- [5] Scott Aaronson and Guy N Rothblum. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 322–333, 2019. DOI: [10.1145/3313276.3316378](https://doi.org/10.1145/3313276.3316378). 21
- [6] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124019, 2019. DOI: [10.1088/1742-5468/ab3988](https://doi.org/10.1088/1742-5468/ab3988). 4
- [7] Dana Angluin and Philip Laird. Learning from noisy examples. *Machine Learning*, 2(4):343–370, 1988. DOI: [10.1023/a:1022873112823](https://doi.org/10.1023/a:1022873112823). 8
- [8] Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. *ACM SIGACT News*, 48(2):41–67, 2017. DOI: [10.1145/3106700.3106710](https://doi.org/10.1145/3106700.3106710). 5
- [9] Srinivasan Arunachalam and Ronald de Wolf. Optimal quantum sample complexity of learning algorithms. *The Journal of Machine Learning Research*, 19(1):2879–2878, 2018. 5
- [10] Srinivasan Arunachalam, Alex B Grilo, and Henry Yuen. Quantum statistical query learning. *arXiv preprint arXiv:2002.08240*, 2020. 5
- [11] Srinivasan Arunachalam, Yihui Quek, and John Smolin. Private learning implies quantum stability, 2021. 4, 5, 20
- [12] Javed A Aslam and Scott E Decatur. Specification and simulation of statistical query algorithms for efficiency and noise tolerance. *Journal of Computer and System Sciences*, 56(2):191–208, 1998. DOI: [10.1145/225298.225351](https://doi.org/10.1145/225298.225351). 8
- [13] Maria-Florina Balcan. Differential privacy and statistical query learning, 2015. 20
- [14] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on computing*, 26(5):1411–1473, 1997. DOI: [10.1137/s0097539796300921](https://doi.org/10.1137/s0097539796300921). 8
- [15] Avrim Blum, Merrick Furst, Jeffrey Jackson, Michael Kearns, Yishay Mansour, and Steven Rudich. Weakly learning DNF and characterizing statistical query learning using Fourier analysis. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, pages 253–262, 1994. DOI: [10.1145/195058.195147](https://doi.org/10.1145/195058.195147). 3
- [16] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)*, 50(4):506–519, 2003. DOI: [10.1145/792538.792543](https://doi.org/10.1145/792538.792543). 7, 22
- [17] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: The SuLQ framework. In *Proceedings of the Twenty-Fourth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 128–138, 2005. DOI: [10.1145/1065167.1065184](https://doi.org/10.1145/1065167.1065184). 3, 20
- [18] M. Bun, R. Livni, and S. Moran. An equivalence between private classification and online prediction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 389–402, Los Alamitos, CA, USA, November 2020. IEEE Computer Society. DOI: [10.1109/focs46700.2020.00044](https://doi.org/10.1109/focs46700.2020.00044). 5, 20
- [19] Hao-Chung Cheng, Min-Hsiu Hsieh, and Ping-Cheng Yeh. The learnability of unknown quantum measurements, 2015. 4
- [20] Andrew W Cross, Graeme Smith, and John A Smolin. Quantum learning robust against noise. *Physical Review A*, 92(1):012327, 2015. DOI: [10.1103/physreva.92.012327](https://doi.org/10.1103/physreva.92.012327). 5, 8

- [21] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014. DOI: [10.1561/04000000042](https://doi.org/10.1561/04000000042). 3, 20
- [22] Vitaly Feldman. A complete characterization of statistical query learning with applications to evolvability. *Journal of Computer and System Sciences*, 78(5):1444–1459, 2012. DOI: [10.1109/focs.2009.35](https://doi.org/10.1109/focs.2009.35). 3
- [23] Vitaly Feldman. Statistical query learning. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*, pages 2090–2095. Springer New York, New York, NY, 2016. ISBN 978-1-4939-2864-4. DOI: [10.1007/978-1-4939-2864-4_101](https://doi.org/10.1007/978-1-4939-2864-4_101). 2
- [24] Surbhi Goel, Aravind Gollakota, Zhihan Jin, Sushrut Karmalkar, and Adam Klivans. Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. In *International Conference on Machine Learning*, 2020. 4, 6, 7
- [25] Daniel Gottesman. Stabilizer codes and quantum error correction. *arXiv preprint quant-ph/9705052*, 1997. 2
- [26] Daniel Gottesman. The Heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*, 1998. 2
- [27] Alex B Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Physical Review A*, 99(3):032314, 2019. DOI: [10.1103/physreva.99.032314](https://doi.org/10.1103/physreva.99.032314). 5
- [28] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, pages 1–1, 2017. ISSN 1557-9654. DOI: [10.1109/tit.2017.2719044](https://doi.org/10.1109/tit.2017.2719044). 2, 4
- [29] M. Hamermesh. *Group Theory and Its Application to Physical Problems*. Addison Wesley Series in Physics. Dover Publications, 1989. ISBN 978-0-486-66181-0. DOI: [10.1119/1.1941790](https://doi.org/10.1119/1.1941790). 11
- [30] Lisa Hellerstein and Rocco A Servedio. On pac learning algorithms for rich boolean function classes. *Theoretical Computer Science*, 384(1):66–76, 2007. DOI: [10.1007/11750321_42](https://doi.org/10.1007/11750321_42). 2
- [31] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011. DOI: [10.1137/090756090](https://doi.org/10.1137/090756090). 20
- [32] Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998. DOI: [10.1145/293347.293351](https://doi.org/10.1145/293347.293351). 2, 8, 9
- [33] Michael Kearns and Ming Li. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 22(4):807–837, 1993. DOI: [10.1137/0222052](https://doi.org/10.1137/0222052). 8
- [34] Ching-Yi Lai and Hao-Chung Cheng. Learning quantum circuits of some T gates, 2021. 4
- [35] Ashley Montanaro. Learning stabilizer states by Bell sampling, 2017. 4
- [36] Ryan O’Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pages 899–912, 2016. DOI: [10.1145/2897518.2897544](https://doi.org/10.1145/2897518.2897544). 2, 4
- [37] Ryan O’Donnell and John Wright. Efficient quantum tomography II. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 962–974, 2017. DOI: [10.1145/3055399.3055454](https://doi.org/10.1145/3055399.3055454). 4

- [38] Krzysztof Pietrzak. Cryptography from learning parity with noise. In *International Conference on Current Trends in Theory and Practice of Computer Science*, pages 99–114. Springer, 2012. DOI: [10.1007/978-3-642-27660-6_9](https://doi.org/10.1007/978-3-642-27660-6_9). 7
- [39] Patrick Rall, Daniel Liang, Jeremy Cook, and William Kretschmer. Simulation of qubit quantum circuits via Pauli propagation. *Physical Review A*, 99(6), June 2019. ISSN 2469-9934. DOI: [10.1103/physreva.99.062337](https://doi.org/10.1103/physreva.99.062337). 10
- [40] Lev Reyzin. Statistical queries and statistical algorithms: Foundations and applications. *arXiv preprint arXiv:2004.00557*, 2020. 2, 3
- [41] Andrea Rocchetto. Stabiliser states are efficiently PAC-learnable. *Quantum Information & Computation*, 18(7-8):541–552, 2018. DOI: [10.26421/qic18.7-8-1](https://doi.org/10.26421/qic18.7-8-1). 1, 2, 3, 4, 12, 21
- [42] Andrea Rocchetto, Scott Aaronson, Simone Severini, Gonzalo Carvacho, Davide Poderini, Iris Agresti, Marco Bentivegna, and Fabio Sciarrino. Experimental learning of quantum states. *Science advances*, 5(3):eaau1946, 2019. DOI: [10.1126/sciadv.aau1946](https://doi.org/10.1126/sciadv.aau1946). 4
- [43] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984. DOI: [10.1145/1968.1972](https://doi.org/10.1145/1968.1972). 1
- [44] LG Valiant. Learning disjunction of conjunctions. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence-Volume 1*, pages 560–566, 1985. 8
- [45] Mithuna Yoganathan. A condition under which classical simulability implies efficient state learnability, 2019. 4, 21